Can a spiking neural network predict unknown learning histories solely on a snapshot of binary choice data?

Anna Plessas

2023

A thesis submitted to Auckland University of Technology
in fulfilment of the requirements for the degree of
Doctor of Philosophy

Faculty of Health and Environmental Sciences
Psychology and Neuroscience Department

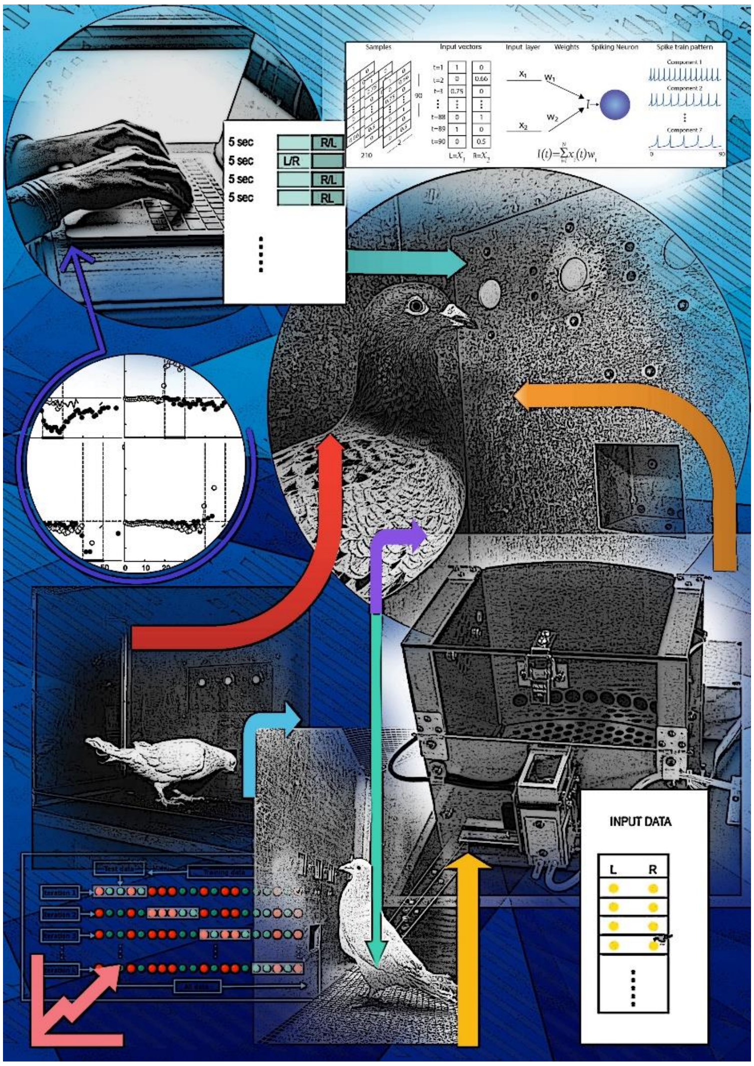

Abstract

Behaviour is influenced by past experiences, allowing us to make predictions based on current behavioural patterns. The aim of this doctoral project was to predict (more precisely, to retrodict) unknown learning histories based on small-sized datasets derived from binary choices made by pigeons. These choice behaviour datasets were extracted as discrete events from published research papers and each data included a 5-second window of observable behaviour after the delivery of a reinforcer. A spiking neural network (SNN) with a single Leaky Fire-and-Integrate neuron was developed to process these data and generate retrodictions of learning histories.

Experiment 1 showed that retrodictions can be made by approaching behavioural data differently, without manipulating the reinforcer-behaviour relationship. Pigeons' binary choices provided sufficient information to the SNN model, which successfully matched the manual analysis of their actual choice behaviour patterns. The SNN's ability to make retrodictions relied solely on the pigeons' unique response patterns and the SNN's learning capacity. Retrodictions were successful even when the SNN was tested with new samples of various sizes from the same datasets. Thus, the SNN model demonstrated its capability to learn and make accurate predictions from behavioural data.

In Experiment 2, the effectiveness of the SNN was validated by comparing its performance with that of other artificial neural networks. Three deep learning models were developed. The retrodictions made by these models were then compared to the performance of the SNN. The results showed that all models were able to accurately retrodict the pigeons' learning history. However, when additional performance measures such as F1 and precision were taken into account, the SNN outperformed all other deep neural network (DNN) models.

Experiment 3 showcased the SNN's capability to work with novel small-sized datasets consisting of choice behaviours of other pigeons, both individuals and groups, who had slightly different learning histories. By recalculating the SNN's firing rates in a personalised manner, better predictive performance was achieved compared to conventional approaches used in generalisation tests, despite variations in the pigeons' learning histories. To confirm the method's reliability, Experiment 4 involved retraining the SNN with new datasets by using two transfer-learning techniques (fine-tuning and feature extraction) and then testing it on small new datasets. Both strategies yielded robust retrodictions, demonstrating the advantages of the applied methodologies.

This thesis may become a valuable addition to the behaviour-analytic toolbox by providing a tool capable of retrodicting unknown learning histories from limited behavioural datasets containing a small window of binary choices. Collectively, the experiments demonstrate that the SNN is an effective tool for understanding the connection between learning histories and behaviour in behavioural research. The SNN exhibited adaptability and responsiveness to relatively small amounts of data from observable behaviour and produced retrodictions, thus demonstrating its potential to replace labour-intensive manipulations of the reinforcer-behaviour relationship or lengthy common training procedures. It makes personalised predictions possible and facilitates the study of differences in individual learning patterns shedding light on the relation between learning history and behaviour. By retrodicting learning histories, this work establishes a foundation for exploring the use of new training methodologies, using optimal training conditions tailored to individual organisms and specific learning tasks.

Contents

List of Figures

List of Tables

List of Appendices

Attestation of Authorship

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person (except where explicitly defined in the acknowledgements), nor material which to a substantial extent has been submitted for the award of any other degree or diploma of a university or other institution of higher learning.

Date: 29-05-23

Candidate contributions to co-authored papers

Chapter publication reference	Author %
Chapter 4. Plessas, Anna, Josafath I. Espinosa-Ramos, Dave Parry, Sarah Cowie, and Jason Landon. 2022. “Machine learning with a snapshot of data: spiking neural network ‘predicts’ reinforcement histories of pigeons’ choice behaviour.” Journal of the Experimental Analysis of Behaviour, April. https://doi.org/10.1002/JEAB.759.	AP 80% J ER 8% DP 2% SC 5% JL 5%
Chapter 5. Anna Plessas, Joshua Bensemann, Josafath I. Espinosa-Ramos, Sarah Cowie, Jason Landon, Dave Parry. "Performance comparison of spiking neural networks to other neural networks in predicting reinforcement history of pigeons' binary choices", 21 March 2023 submitted for review at Neural Computing and Applications	AP 80% JB 10% JER 1% SC 1% JL 4% DP 4%
Chapter 6. Anna Plessas, Josafath I. Espinosa-Ramos, Sarah Cowie, Jason Landon, Dave Parry. “A spiking neural network to personalise learning histories for pigeons’ choice-making”, to be submitted for review	AP 80% JER 10% SC 2% JL 2% DP 6%
Chapter 7. Anna Plessas, Josafath I. Espinosa-Ramos, Sarah Cowie, Jason Landon, Dave Parry. “Where do pigeons’ choices come from? Transfer learning in spiking neural networks for animal choices in behavioural research”, 19 March 2023 submitted for review at Neural Computing and Applications	AP 80% JER 8% SC 2% JL 5% DP 5%

We, the undersigned, hereby agree to the percentages of participation to the chapters identified above

Jason Landon,
Associate Professor

Dave Parry, Professor

Sarah Cowie,
Associate Professor

Josafath I. Espinosa-Ramos, PhD

Joshua Bensemann, PhD

Acknowledgements

I am deeply grateful for the trust, kindness, and nurturing that my supervisors have consistently gifted me throughout this incredible journey. Thank you for your unwavering dedication and belief in my abilities. Your contributions have been invaluable, and I have been extremely fortunate to have the opportunity to work under your guidance.

I extend my sincerest appreciation to Associate Professor Jason Landon, whose unwavering support and encouragement have pushed me to be more explorative and to challenge my own ideas. Your honest and constructive feedback has been instrumental in expanding my knowledge and honing my ability to effectively communicate my ideas. I would also like to express my heartfelt thanks to Professor Dave Parry for instilling confidence in me. Your belief in my abilities and your guidance in finding novel solutions have been instrumental to the success of this research. I am indebted to Associate Professor Sarah Cowie for her generous sharing of her expertise in choice research and for guiding me to delve deeper into the realm of behavioural science. The opportunity to visit the Pigeon lab and glean valuable insights into my data has been truly enlightening. Your guidance has contributed significantly to my growth and knowledge in this field.

To all of you, I extend my gratitude for treating me with respect, attentively listening to my evolving ideas, and steering me in the right direction. I cannot envision a more inspiring team of supervisors who have consistently demonstrated how a cohesive and supportive team should work.

I would like to express my gratitude to Dr. Israel Ramos Espinosa for his invaluable teachings on spiking neural networks and for guiding me to understand how to evaluate machine learning research. Your expertise and knowledge in these areas

have been instrumental to expanding my understanding and skills. I would also like to extend my thanks to Dr. Joshua Bensemann for being open to collaboration and for sharing the same passion for leveraging behaviour analysis through the use of automated systems. Your willingness to work with me has been inspiring, and our collaboration has allowed for significant progress in this research area. Thank you Nektarios Papakonstaninopoulos for teaching me Python and the importance of learning how to navigate through my datasets. Thank you, Sylvia Kaprou, for accepting my scribbles and turning them into art.

Lastly, I what to thank my husband and companion in life, Ignatios Kafantaris, who contributed to the idea for this research, for his unwavering support and constant encouragement. Your belief in me and your continuous push to achieve more have been the pillars of my strength throughout this journey. I want to thank my daughters Amira and Ariadne for putting up with mama and loving me despite them hearing "*I'm almost done with my writing*" a hundred times. Thank you, my little Ariadne, for your heartfelt Mother's Day message (verbatim) "*She inspires me by not giving up, commits to something and want stop and finishing everything on time….and I will always listen to you, love you and be kind*". Thank you, dad and mum, Leonidas and Ioanna Plessas; though it was hard for you to understand why I was stubbornly doing all this, you were always supportive and helped us to rear our children. Thank you for being my pillars of strength and for instilling in me the values of hard work, determination, and belief in myself. I am forever grateful for your love, guidance, and unwavering support. Σὺν Θεῷ τα καταφέραμε. Μέγας εἶ, Κύριε, καὶ θαυμαστὰ τὰ ἔργα Σου. Εὐλόγησον.

Ethics Approval

All data were extracted from experiments conducted at the experimental analysis of behaviour lab of the University of Auckland, New Zealand using operant procedures to study choice behaviour. No additional ethics approval is necessary for the use of this data. All data3 have been provided by the original investigators and written permission has been obtained (Appendix A)

A definition of prediction

According to the Cambridge English dictionary, prediction is “to say that an action or event will happen in the future.” However, when this term is used in artificial intelligence or psychology, the meaning will not exactly fit with this definition.

The ultimate goal in machine learning (ML) is to automate tasks by creating models with algorithms for processing data (more on this in Chapter 2). Algorithms train a computer model on available historical data, so the model can predict future, or else, unseen data. Prediction relates to the output of an algorithm, and the goal is to produce an accurate output/prediction. An accurate prediction is validated when the trained model is applied to new data and is able to generate a new accurate outcome. The purpose of the predictive outcome is to find an answer to an applied problem by detecting patterns in the dataset. For example, an interesting problem may be for an ML algorithm trained on historical data of a patient’s therapy to predict if an alternative therapy can offer better outcomes. The model is not constrained by the interpretability of the outcome as long as the outcome is valid.

In machine learning papers, we observe the term ‘classification’ being used instead of ‘prediction’, which may be misleading to the naïve audience. Classification refers to a specific predictive problem, where the goal is for the algorithm to generate a categorical outcome, in other words to classify the data according to their pattern similarities. Accuracy of classification is what makes the outcome predictive. Therefore, classification and prediction are used interchangeably.

In psychology, the ultimate goal is to understand human behaviour. Traditionally, this is achieved by explaining a behaviour by manipulating independent variables that can be causally related to the behaviour; as a result of such findings, a behaviour that has not been yet observed may be predicted. To predict would imply that it is possible to control the behaviour. Therefore, in practice, understanding and predicting are intertwined.

In this doctoral thesis, prediction and explanation of an event are treated as two compatible approaches as research recommends (Killeen, 2019; Yarkoni & Westfall 2017). Furthermore, in this research project, the goal was to 'predict' the learning history of a current observed behaviour (more in Chapter 1). However, it would have been more accurate to use the term retrodiction. Retrodiction signifies that the unseen event is in the past rather than it is one that has not occurred yet. However, in machine learning, prediction implies that a shape computed by an algorithm fits as close as possible to the data. Often terms borrowed from neuroscience can be confusing when their use serves a different function. For example, when computer science borrows from neuroscience the noun 'intelligence' to construct the term 'Artificial Intelligence', it can be mistaken as machines having cognition, although AI is only a set of mathematical computations that execute specific tasks.

To keep consistent with the ML terminology, the term prediction is used interchangeably with retrodiction to indicate that the generated outcome by the machine is an unknown (i.e., unseen) that took place in the past.

Chapter 1 Introduction

This doctoral thesis aspires to contribute to developing a method of analysis of behavioural spatiotemporal data in the Experimental Analysis of Behaviour (EAB) to identify learning histories. Behaviour is generally understood in relation to past experience; therefore, current patterns in behaviour should allow retrodictions of past experiences. This would be interesting both from a theoretical perspective and because it has implications for practice.

The vertiginous development of successful artificial intelligence (AI) applications in the last decade or so has triggered many scientists, including psychologists, to revisit the utility of prediction. For an experiment, a hypothesis, control and prediction are the essential ingredients. Firstly, independent variables (ingredients that remain unchanged) are tightly controlled in order to test and evaluate the observations with the results. Secondly, prediction can be inferred based on the results that validate or disprove the hypothesis. Yet, control in experiments can be hard (Torday & Baluška, 2019), making prediction a difficult goal to achieve.

Using AI for prediction can complement traditional experimental approaches if suitable datasets are available. The Experimental analysis of behaviour (EAB) is a branch of psychology that studies environment-behaviour interactions and has generated an abundance of such data.

Huge efforts of EAB researchers have allowed a level of prediction possible, but there is a possibility that emerging technologies could facilitate this further. The existing technology used in EAB experiments is both expensive and time-consuming, which could disincentivise most researchers to work on questions like: What if the behaviour of an organism is known, but the learning history is unavailable to us? How can we ‘predict’ this information in order to optimise new learning experiences for that individual? Here, learning is defined as the interactions between organism and environment that shape the organism’s behavioural repertoire.

The emergence of AI is an incalculable leap in productivity, and today it presents an opportunity to rethink the attainability of finding answers to questions that were not posed hitherto or did not get traction. Therefore, AI opens an avenue to pose alternative questions to solve existing problems, even as traditional research continues.

Two overarching questions are addressed in this thesis:

1. If only a snapshot of choice behaviour data is available, can we ‘predict’ the learning history – the reinforcer arrangement used in training conditions– that led to the current choice behaviour observed by us?
2. Suppose prediction modelling (an AI-type algorithm) helps us to answer question #1. What happens when we want to predict the learning history of other participants or learning histories that differ from the ones used for the development of the AI model? Will the predictive outcome be accurate?

The novelty of this approach and the limited research in the area, require some background on the behavioural research behind these two questions (Chapter 1 – this chapter) and on the AI technology that can assist us to achieve the desired prediction (Chapter 2). The research questions, the rationale and the significance of this thesis are explained in more depth in Chapter 3.

This chapter will present the science and research behind the datasets used in this doctoral thesis. First, a definition of the science of behaviour analysis is provided, and how behaviour is viewed in context. Then some history and basic principles for quantifying observable behaviour are followed by measures of choice behaviour, constituting this thesis's datasets. Lastly, the experimental procedures used to generate choice behaviour and the variables affecting response patterns (including learning history) are presented. Throughout Chapter 1, the emphasis is on prediction, which is the overarching theme of this thesis.

1.1 Behaviour Analysis defined

Behaviour analysis is a branch of psychology that studies learnt behaviour in its own right, not as an index of other underlying events. Behaviour analysis is not premised on dualism: this means that behaviour is not just a datapoint of an underlying mental state or cognitive system (Schlinger, 2018; B. F. Skinner, 1989).

Behaviour refers to anything an organism does, and it is studied in relation to environmental controlling variables. There is constant interaction between organism and environment, whereby environmental changes are followed by the organism responding to them, thereby changing the environment, and so on. Every event has a feedback function that can be described mathematically or graphically to demonstrate how the environment presents contingencies as a function of an organism's behaviour (Baum, 1973). Hence, in behaviour analysis, there is a strong focus on environment-behaviour interactions to explain how behaviour can change, constituting new learning, to meet the demands of an ever-changing environment (Pierce & Cheney, 2004). To develop an understanding of how learning takes place, experimental research focuses on control and systematic observations of behavioural changes while the dynamic interaction between environment and behaviour evolves. Behaviour analysts use the term *adaptive learning* to refer to behaviour that is acquired, maintained and changed through an organism's life. From early on, behaviour analysts theorised that responses can also have an altering effect on the organism’s current state (e.g., brain functions) when learnt behaviours are changed (Baer, 1976). Research has suggested that the study of adaptive learning can provide a framework to investigate the neural bases of behaviour, as behaviour is characterised by variability analogous to neural plasticity and genetic variation (Schlinger, 2015).

There are two primary areas of behaviour analysis: experimental and applied. Experimental analysis of behaviour (EAB) focuses on the research of adaptive behaviour in laboratories to formulate an empirical and theoretical understanding of learning processes for human and non-human organisms (Skinner, 1966b). Applied behaviour analysis (ABA) focuses on the application of this knowledge to real-life situations. ABA research is conducted in a similar manner, with a focus on understanding what controls behaviour. However, the questions asked by ABA are applied and tackle a socially significant behaviour to improve a person's quality of life (Baer et al., 1968).

Therefore, the term *behaviour analysis* implies that behaviours are studied using specific methods and techniques, followed by an analysis focused on the behaviour's interplay with the environment. This thesis refers to *behaviour* under the theoretical framework of behaviour analysis.

1.2 The science of behaviour

1.2.1 Introduction to prediction

Prediction has been a fundamental standard in psychology since the beginning of the last century (Sarbin, 1944). Predicting a living organism's behaviour requires the behaviour analyst to develop research methodologies that make possible the prediction of future behaviour based on current behaviour. In laboratories, prediction of future responses is inferred by strictly controlling the environmental conditions. Despite the precision in executing experiments to quantify relations, the elegant experimental designs, and the powerful results EAB has accumulated over the years, the body of research has produced more descriptive rather than predictive accounts of learning mechanisms.

Prediction and control became the fundamental aim in EAB studies of the *functional relations* between environment and behaviour. By *function* we refer to the effect produced either by the behaviour or by controlling variables in the environment (Pierce & Cheney, 2004). Hence, a functional relationship exists when the manipulation and variability of one (environmental conditions) have an observable effect on the other (behaviour). The functional relationship is inferred, based on observations of behaviour that is shaped by the environment and individual differences (e.g., genetic factors). A classical example of how EAB builds its learning theory is by observing pigeons pressing on keys and recording changes in their pecking behaviour in response to specific environmental arrangements, rather than making deductive inferences on pecking (Chiesa, 1994).

EAB is methodologically very different to other streams of psychology. It is based on the analysis of data from within-subject designs with small groups of participants, rather than from between-subject designs and big groups. Under strictly controlled environments, behaviour changes are measured and analysed to attempt prediction, without the use of hypo-deductive methods. Following is an introduction to how EAB has come to understand adaptive learning.

1.2.2 The historical evolution of ‘Behaviour-Reinforcement’

John B. Watson, a pioneer psychologist of his time, questioned how psychological research was conducted and proposed that research methods respond to natural science in order for psychology to be the science of observable behaviour (Watson, 1914). He was one of the first to argue that experimentation in psychology should divert from subjectivity and focus on control and prediction, like any other natural science. Therefore, observable phenomena should be the focus of study and not conscious/unconscious experiences.

Around the same time, Edward Thorndike, another influential psychologist, published a series of experiments to demonstrate that studying observable behaviour is possible, and he developed the *Law of Effect* (Thorndike, 1911). He demonstrated that if a current behaviour originates from a similar environment to a past one, where the behaviour produced a satisfying effect, it will likely be the same, due to that past learning experience. Thorndike demonstrated that control and prediction are possible in the study of behaviour, and he argued that past environments likely increase the strength of the same behaviour in similar future environments.

Later, for B.F. Skinner, it very quickly became apparent that behaviour is probabilistic even for well-defined patterns of behaviour and it is affected by how the behaviour and the environment interrelate. The emergence of an observable behaviour was understood by Skinner as a function of complex and subtle phenomena, and he directed his research to environmental events and responses (*The Experimental Analysis of Behavior | American Scientist*, n.d.). The occurrence (or non-occurrence) of a behaviour was studied under tight experimental conditions, to control when the behaviour recurs. Skinner refined Thorndike's work by showing it is the *reinforcer* –a controlling consequence that is specifically delivered following a behaviour–) that increases the future probability of the behaviour to recur (Skinner, 1938). Therefore, the experimental environment was engineered by manipulating *reinforcer arrangements* and observing changes in the organism's behaviour (Ferster & Skinner, 1957; Skinner, 1966a; Staddon & Cerutti, 2003; Thorndike, 1911). Skinner's experiments with pigeons led to the publication of his seminal work 'Schedules of reinforcements', a set of defined rules that dictate how reinforcers can be manipulated to strengthen a specific behaviour’s probability to occur in the future (Ferster & Skinner, 1957). Skinner suggested that knowing the causes of behaviour, was absolutely necessary for prediction and control (Overskeid, 2018). According to (Skinner, 1981), the 'selection by consequences' is akin to natural selection and the organism's actions cannot be credited to an autonomous agent with reasons for acting. Therefore, learning is an adaptive process to the changing environment that we can observe through the changes in a living organism’s behaviour.

Taken in a simplistic way, anyone may try to apply these predefined rules; a teacher, a parent, a therapist or a researcher may arrange (or withdraw) a reward, after a behaviour has occurred, based on timing or frequency rules, with the expectation that this experience will alter the individual’s future responses. However, anyone who has attempted to change a behaviour by exclusively relying on a causal response-reinforcer relationship, will have quickly realised that change is not so straightforward. Providing consequences just by setting a predefined rule, predictably leads sometimes to false-positive outcomes. Though Skinner’s work was incredibly important to the field, a series of contemporary EAB papers has demonstrated that Skinner’s conclusions hold true only when the response is likely to produce a future reinforcer (e.g., Cowie & Davison, 2020b). Thus, past experiences certainly play a role in shaping future behaviour, but not an exclusive one. Before presenting current research and how the learning history contributes to predictions, research behind choice behaviour and its contributions in understanding adaptive learning will be presented. This is necessary for the reader to understand why choice behaviour datasets were selected for this project.

1.2.3 Why study choice behaviour?

B. F. Skinner's experimental analysis and the quantification of behaviour for measurement and analysis, led to important contributions during the 1960s. The experimental method involved breaking down these complex relations through highly controlled procedures using a device called the *operant chamber* (Figure 1). An operant chamber is a small box designed for a non-human organism that presses a lever or pecks a key when lights illuminate. A cup is connected to an electronic feeder for food to come out after a certain number of behaviours occurs, or after a set time following the 'emission of target behaviours'. The basic procedure commences by showing the availability of two alternatives for the participant to choose from; each alternative key has a reinforcer arranged for different contingencies, namely the dependent relation between behaviour and reinforcer arrangements. Responses on one key in relation to responses to the alternative one are measured. *Concurrent schedule procedures* provide a means to make these arrangements explicit and record the behaviour for analysis (more in Chapter 4). The focus is on factors affecting behaviour over time.

Figure 1

Pigeons experimental chamber in behavioural research laboratories

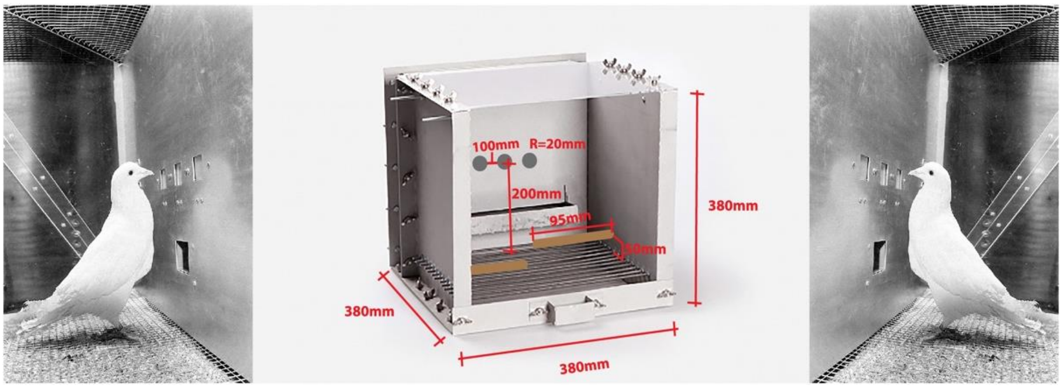

Artwork was designed for experiment in Chapter 7

Choice behaviour was studied as the main paradigm of behaviour because it involves the distribution of responses to alternative keys and the participant must choose from behaviours that exist in their own repertoire (Herrnstein, 1970). Thus, all behaviour is choice behaviour, as it involves doing one thing or the other. Herrnstein (1970) explained that even in a simple environment, like the operant chamber, where only one-key is available and the reinforcer is arranged only for that key, the participant has the option to choose to respond to that key or to other reinforcers available. For example, it may walk around the chamber, look outside the cage, scratch, and so on. In this lies the rationale for the extensive research in choice behaviour and the variables that control it.

In summary, Herrnstein (1970) explained that every response at any given time is a choice behaviour, and under this conceptual framework, research has focused on understanding the interplay between competing reinforcers and the power of switching contingencies from favouring one choice response such as a pigeon pecking on the left key, to favouring an alternative one by pecking on the right. Choice behaviour was recorded as a datum of response rate. Herrnstein's work influenced the study of control, thus prediction, of learnt behaviour by presenting a linear relationship between the increase in responses and the relative rate of reinforcers for that behaviour (Plaud, 1992).

The Harvard pigeon lab played a significant role in developing Behaviour Analysis as a science in two significant periods. The time where Skinner and his students navigated the science of behaviour as a natural science and established concepts such as the schedules of reinforcement was the first period. The second period was ushered in by Herrnstein where choice behaviour was quantified and new notions were introduced, such as the relative response rate, choice behaviour, and the Matching Law (Baum, 2002).

1.2.4 Quantitative models of choice behaviour

The purpose of this section is not to provide an extensive overview of all the experiments that used quantified models to demonstrate the interplay between behaviour and environmental conditions. This short introduction aims to show how control was represented, data was collected and analysed. The great pursuit of EAB experimenters was to achieve prediction.

Herrnstein (1961) conducted a study using pigeons as participants and measured their response rates when specific reinforcer arrangements (i.e., *variable-interval concurrent schedules of reinforcement*) were made for each key (Herrnstein, 1961). In Herrnstein's study, the reinforcer arrangement was held constant across several trials for each key, and when the pigeons switched from one key to another, the reinforcer was withheld for a very short period for the first response to occur (i.e., 1.5 second of *changeover delay*). This seminal work demonstrated that relative response rates equalled or matched the reinforcer arrangement under these conditions and it generated the assumption that behaviour is not only related to its immediate effects but also to the consequences of alternative behaviours. Thereby, a mathematical equation denoted as the *matching law* was formulated.

Put simply, behaviour matches reinforcers, and the matching function was formulated as

$$\frac{B1}{B1+B2}=\frac{R1}{R1+R2} \qquad (1)$$

where B1 and B2 are the numbers of emitted alternative behaviours, and R1 and R2 are the numbers of reinforcers delivered for the respective alternative behaviours. Perfect matching would predict that changes in reinforcer delivery indicate with 100% accuracy, the changes in relative behaviour in those conditions. It also requires that the organism distribute its behaviour equally between all the reinforcers obtained. The matching law was studied extensively in the 1960s and was applied across non-human and human organisms, for all sorts of reinforcer arrangements and responses, and systematically showed that behaviour was distributed according to the relative distribution of reinforcers. However, perfect matching was not a common result. Extensive research showed the relationship was non-linear, the distribution of responses was typically less extreme than the distribution of reinforcers, a finding commonly referred to as undermatching when these are measured as proportions. When the logarithm of response ratios was plotted as a function of the logarithm the reinforcer ratio, the relationship is a straight line, but the slope was generally less than 1.0, as perfect matching would demand. Further, the y-intercept is not always zero (Hollard & Davison, 1971; Taylor & Davison, 1983). These results indicate some sort of bias, sometimes due to experimental factors and sometimes for other reasons related to individual differences. Nevertheless, it introduced the notion that occurrences of behaviour are relative; and response patterns are a function of all reinforcers available as consequences. Therefore, in simple terms, a therapist will have more success when a consequence (a reinforcer) is given in a context where other reinforcers are also available.

The common results of undermatching and bias (Baum, 1973, 1974) were accounted for in a modified version of the matching law known as the 'generalised matching law' (GML) (Baum, 1974) in a logarithmic form, as

$$\log\left(\frac{B1}{B2}\right) = a\log\left(\frac{R1}{R2}\right) + logc \qquad (2)$$

where parameter a is sensitivity to the reinforcement, the slope of the best-fitted line, and measures the degree to which the response ratio changes with variations in the reinforcer ratio (Lobb & Davison, 1975). Log c represents inherent bias, or how much preference the organism has over one response that cannot be accounted for by the reinforcer delivery and is explainable by other factors, such as the physical characteristics of the organism (Baum, 1974). Strict matching implies that sensitivity to reinforcement (a) is 1.0 and that the ratio of behaviours must be sensitive to the ratio of reinforcers. Experiments using concurrent schedule procedures have shown that sensitivity to reinforcement is often less than 1.0 (e.g., Baum & William, 1979). Consequently, the GML model provides an excellent means of describing behaviour patterns, given that the relative rate of the reinforcer can describe the behaviour. However, beyond predicting a linear relationship between the logarithms of response and reinforcer ratios, the GML is a descriptive model. For a model to be predictive it should be able to provide information on why behaviour continuously does not fit with the model. One possible explanation of why these models account only for descriptions and not predictions, is that prediction is based on the available data only and not on the basis of logistic functions (Killeen, 2015). Also, the GML models make the implicit assumption that choice behaviour is homogenous across time, following reinforcer arrangements and the control is apparent when large aggregations are used in analysis. However, this approach may not apply for data where the response patterns are not equal or when the value of reinforcers changes between the two alternatives.

1.2.5 What changed? Variations in environment affecting choice behaviour

The vast majority of experiments were set up to observe behaviour in what is called a *steady-state procedure.* That meant that the reinforcer conditions remained the same with the passage of time, up to 30 daily sessions and only changed according to set criteria (Davison & McCarthy, 1988; Killeen, 1978). Researchers allowed an adequate number of sessions for choice behaviour to reach a stability criterion, and the data from these stable periods were subjected to analysis. However, this approach alone can be time-consuming, labour-intensive, and resource intensive. The data analysis, however, typically did not examine data in a more fine-grained way, for example response-by-response, or as choice behaviour was transitioning between stable states. In recent years, the attention of some researchers turned to behaviour in changing conditions, and examining whether more information could be garnered about the mechanisms underlying the pigeons' choice making.

Researchers started to investigate how different variations in aspects of the environment led to different levels of behavioural control. This environmental variation was represented as changing the reinforcer arrangements when using concurrent schedule procedures. This change was done between sessions, changing one aspect of the environment within the session, or having multiple changes within the session. For example, research changed the reinforcer ratios presented in a concurrent schedule between sessions, to examine how control in a previous session can control choice behaviour in the present session (Hunter & Davison, 1985; Schofield & Davison, 1997). One important implication was that data could potentially be collected in shorter time frames, which would reduce costs while still providing valuable insights into how choice behaviour changes, as in stable conditions. They showed that the previous reinforcer ratio had a large effect on choice making in the first session and the control disappeared in three to six sessions (Hunter & Davison, 1985). They also showed that control is stronger when choice behaviour occurs proximally to reinforcer delivery (i.e., known as localised control, as the local choice matches the reinforcer ratio)(Schofield & Davison, 1997). Other studies varied the reinforcer ratio unexpectedly within a session in various ways. For example, Mazur and his colleagues conducted several studies, and they concluded overall, that when the environment changed frequently, the behaviour changed to a greater or lesser extent depending on the environmental context in which a reinforcer was delivered. Also, the localised effect of reinforcers reduces with the passage of time, while previous reinforcer ratios affect the control of choice behaviour (e.g., Mazur, 1997). These and other studies concluded that control of choice behaviour can be affected in present time by the immediate environment but also from the overall environment arrangement (i.e., the frequency of alternations).

Other experiments changed the reinforcer ratios multiple times within sessions. This gave rise to a procedure that has been used in recent years, which has investigated choice behaviour in *frequently changing environments*. Resembling what may occur in a natural environment, the research is conducted under variable reinforcer arrangements and overall reinforcer rates, in order to illustrate how this variability may affect the acquisition and distribution of behaviour in a concurrent schedule procedure. A seminal work was presented by Davison and Baum (2000) who adapted a procedure introduced by Belke and Heyman (1994) where seven different reinforcer ratios were arranged as components in one single session in a random order. This procedure gave a means to examine choice behaviour under conditions that changed rapidly, therefore unpredictably to the participant. They investigate how choice behaviour can change, based on the frequency of environmental change and the overall reinforcer rate (Davison & Baum, 2000). They observed that the participants learnt to shift their choice behaviour quickly in relation to the rapid changes of reinforcer ratio. The first four to six reinforcers in a component following a

10-s blackout, had a large effect on choice behaviour. The overall reinforcer ratio was also manipulated, demonstrating that behaviour change was faster when the overall rate of reinforcement was higher. The authors also investigated the reinforcer effects in context. They concluded that the sequence in reinforcers influenced choice behaviour and the most recently obtained reinforcer had the largest effect on current preference and that previous reinforcers had a small measurable effect over time (Davison & Baum, 2000). These effects were larger when the reinforcer was given in extreme quantities or ratios respectively (Boutros et al., 2010; Landon et al., 2002). Also, sequences of successive reinforcers obtained from the same alternative had cumulative effects and, when the sequences themselves occurred with greater frequency, their effects were cumulative too (Landon et al., 2003a). The Landon and Davison (2001) procedure, followed Davison and Baum's (2000) work and by using a constant (rather than a variable) overall reinforcer rate and varying the range over which the seven component reinforcer ratios altered. Similar effects were evident, but they were enhanced and reduced as the range of reinforcer ratios was increased (and decreased).

The intention here is not to provide an exhaustive list of all the procedures used and their outcomes. Rather it is to show that through laborious experimentation, different procedures and methods of analysis were adapted to understand that variables controlling choice behaviour, operate at several different levels. At one level, reinforcers affect subsequent behaviour. At another, successive reinforcers have cumulative effects on choice behaviour. At a third level, when sequences occur with higher frequency, these cumulative effects are magnified.

Studies are ongoing and have shown that when the availability of reinforcers (reinforcer ratio) changes across time since some event, choice behaviour tracks these changes (Cowie et al., 2011, 2013, 2014; Cowie & Davison, 2016a; Krägeloh et al., 2005; Miranda-Dukoski et al., 2014). Reinforcers themselves can function as discriminative stimuli and signal information about future reinforcers. New versions of the mathematical models are in the making to better facilitate a closer understanding of the interplay between time, reinforcers, and choice behaviour. In a way (Cowie & Davison, 2016a) the evolution of research demonstrates the limitations of models like the GML, e.g., the predominance of certain factors in experimental design and data interpretation, that do not adequately account for the influence of time as a controlling variable. As a result, these models can sometimes make counterintuitive predictions for specific environmental arrangements.

The choice behaviour data used in this project and the particular setup of those experiments, are presented in more detail in Chapter 4 and Chapter 6.

1.3 Conclusion: Is control necessary for prediction?

Quantitative models of non-human and human choice behaviour need a degree of complexity and flexibility. Many factors can influence choice behaviour such as reinforcer magnitude, ratios, reinforcer probability, individual differences, time elapsed, and so on. These factors can influence the ability of mathematical models to quantify the functional relation between behaviour and environment. The primary goal has been to provide descriptions of how behaviour changes, in relation to the reinforcer arrangement. In fact, the EAB approach (i.e., research in choice behaviour) has yielded fundamental research in mechanisms of adaptive learning as opposed to ways of predicting choice behaviours. What if two behaviours have similar patterns of responding, but have been shaped in different reinforcer arrangements? What can the present inform us about the past and, potentially, the future?

Predicting involves making claims about something that will happen or is not known, often based on information from the past and current behaviour. What if we want to make predictions when information on the environment that shaped the present behaviour is not available or is limited? A predictive mathematical model needs data, and if data are unavailable, then prediction may not be possible or may be inaccurate.

In traditional experimental training procedures, control is fundamental to prediction; if a behaviour is controlled, that means it can be predicted. However, this may not always be the case. For example, it would be an inaccurate to state that when data fit a curve (e.g., matching line), then the curve can predict the data. Just because it is possible to fit existing data to a curve, does not mean that the relationship used to draw the curve, is useful for predicting future relationships. The only claim that can be made is that the curve accounts for the data at hand (Davison & Elliffe, 2009; Killeen, 2019). Killeen (2019) was unambiguous in his argument on Herrnstein’s model: though the model has been replicated numerous times under strict control in identical conditions, when it was tested for its predictive qualities, it failed. Thus, replication alone does not necessarily lead to prediction. Killeen argued that while a model may be successfully replicated in controlled and identical conditions, that does not guarantee its ability to make accurate predictions in different situations or settings. The predictive qualities of a model should be assessed separately, taking into account its performance in diverse and variable environments. He concluded that successful prediction in the EAB field was infrequent.

On one hand, as a minimum requirement, control requires sufficient information to be available, so that the observable phenomenon is controlled by the researcher’s manipulations and not by extraneous variables. On the other hand, prediction does not require all the information to be available, as the task is precisely to acquire the missing information. In a sense, control is about ‘how’ the observed phenomenon has happened in specific environmental conditions, whereas prediction tells us, ‘what’ the specific outcome is. Control aims to understand, and prediction aims to identify what will happen, or what has happened although this information is not available to us: ‘retrodiction’. Therefore, there are two different techniques serving different goals while complementing each other, and both contribute to our understanding adaptive learning processes. Also, control and prediction can help each other in certain situations. For example, what if just a small amount of data is available, and information is missing to carry on a traditional training procedure? Can prediction extract this information to carry on conducting an experiment that needs to be controlled?

Predicting without first controlling a relationship, suggests that a tool capable of extracting information directly from the data, expands the possibilities for analysing behaviour. Therefore, in this thesis, we are implying that control and prediction do not necessarily need one to be derived from the other. Both can serve a different function that is useful for behaviour analysis and offer novel avenues for research and application in the field.

1.4 Why predict the reinforcer history?

Extensive research on choice behaviour has validated previous findings that past and present events are crucial variables that influence the occurrence or non-occurrence of behaviour. Different reinforcer arrangements using concurrent scheduled procedures, have proven to maintain different patterns and rates of choice behaviours. While the delivery of reinforcers remains a significant factor in this relationship, it is no longer considered as central as it was in the theories of Thorndike or Skinner. Instead, reinforcers are recognised as one among many stimuli that influence current behaviour. Consequently, current research is heavily oriented towards understanding in what ways, past and present events, affect behaviour to predict future instances of behaviour.

In particular, early work in the field suggested a direct relationship between future behaviour and its past (Skinner, 1938), but recent research has revealed a more nuanced understanding of behaviour control. It has been recognised that the influence of the environment is not solely retrospective but also includes relational, prospective, and goal-directed characteristics (Cowie, 2018). The learning history of a behaviour is crucial in comprehending its goal-directed nature, as two seemingly similar behaviours can be controlled or affected differently by different aspects of the environment (Bouton & Balleine, 2019). Environmental changes affect how the organism responds, resulting in a continuous cycle of behaviour-environment interaction (Bouton & Balleine, 2019). The ‘functioning’ of reinforcing arrangements depends on factors such as the organism's motivational state, reinforcer context, and the history encoded in the present state, all of which can shape choice behaviour (Killeen & Jacobs, 2016). Therefore, the organism will work towards reinforcers, if it is in the relevant state to do so (Shahan, 2017). For this reason, new research has introduced a different perspective on the role of the environment in controlling behaviour. It suggests that the past, particularly learning histories, act as a building block for the future, and current behaviour is controlled by consequences that are spaced out or temporarily distant from future behaviour (Cowie & Davison, 2020b).

Regardless of the precise impact of past exposure to reinforcing conditions, such as learning history, on an organism's behaviour, and the ongoing debate surrounding the mechanisms of this relationship (e.g., as discussed by Cowie, 2019; Simon et al., 2020), all research acknowledges that past learning and present behaviour, play a significant role in our understanding of adaptive learning. Learning history can provide helpful information as it brings insights into the organism’s previous experience with current behaviour, and their contingencies to understand factors that shaped behaviour across time. This information can help in predicting how an individual may respond to future environmental and contextual factors that have influenced its choice behaviour. Also, it allows researchers and practitioners to tailor their approaches and interventions based on the individual's unique learning history, increasing the effectiveness of choice behaviour change strategies.

Current quantitative models require knowledge of the learning history to understand changes in behaviour. Today new predictions are not attainable if the learning history is not available. The field of artificial intelligence may be the place to look for alternative tools to potentially detect the learning history when this history is unknown to us. The temporal elements in choice behaviours may allow us to derive conclusions about learning that took place in a specific environment.

Answering a novel question such as ‘what is the learning history of the choice behaviour I am observing?’ has both a clinical and a basic research value. Prediction of learning histories will lay foundations for novel approaches to the use of new training methodologies, using optimal training conditions for specific learning and for specific individuals.

Chapter 2 Artificial Intelligence (AI) and Prediction

Chapter 1 presented the science of behaviour analysis and the research behind the datasets used in this doctoral thesis. Modelling choice behaviour in relation to its learning history has presented itself as a nuanced and challenging endeavour, often involving significant resources and complexities. However, it remains crucial to identify the underlying mechanisms of learning patterns in choice-making. Modelling the behaviour of nonhumans and humans has proven to be intricate, as numerous factors come into play, including the influence of reinforcer control and the impact of environmental dynamics on behavioural change (Cowie & Davison, 2016a).

In traditional experimental procedures, control allows prediction. However, In Chapter 1, an alternative approach was proposed, suggesting that control and prediction do not need to be interdependent and can have distinct functions in behaviour analysis. For example, if we are to design optimal environments for learning to occur, prediction (i.e., to identify the unknow learning history) could assist us in reaching this goal by circumventing traditional training procedures. Furthermore, exploring the learning histories of choice behaviours can provide valuable insights into their underlying patterns. Such patterns may be revealed by using alternative tools and asking alternative questions, such as: "What is the learning history of the behaviour I am observing?". This approach expands our understanding beyond the traditional methods and opens new avenues for analysis.

Science relies on fitting models to observational data to extract knowledge from them. Similarly, behaviour analysis relies on strict control of reinforcer manipulanda to measure and analyse behaviour changes. However, this traditional approach can be labour-intensive, requiring specialised expertise and manual analysis that can be costly. The advancement of automated computer models offers alternative methods of conducting science. By automating the learning process, these models have the potential to extract and analyse information automatically, reducing the reliance on manual analysis (Alpaydin, 2020). Also, this technological advancement has the potential to improve efficiency while effectively expanding our knowledge from existing experimental studies. Experimental studies provide the datasets for an artificial model to train on. In this thesis, it is suggested that by embracing alternative perspectives, exploring learning histories, and leveraging automated computer models, we can enhance our understanding of behaviour and advance the field of behaviour analysis.

The novelty of this approach and the limited research in the area requires some background on the behavioural research behind the two questions (Chapter 1) and the AI technology that can assist in achieving the desired prediction (Chapter 2 – this chapter). The research questions, the rationale and the significance of this thesis are explained in more depth in Chapter 3.

This chapter will present how AI has been used in psychology and behaviour analysis, followed by the presentation of the AI tool of choice in this research: a third-generation, brain-inspired model, namely a spiking neural network (SNN) requiring only small datasets. At the end of this chapter, a summary of Chapters 1 and 2 with conclusions is provided.

2.1 AI, ML, ANN, DNN, SNN, DL: What is this all about?

Artificial intelligence (AI) is used to denote the development of computing technologies that are able to perform tasks that only humans would have the ability to complete (*Artificial Intelligence - Reasoning | Britannica*, n.d.). The evolution of artificial intelligence (AI) over the last 15 years, has seen the advancement of algorithms, improvements in storing data, and the need for minimal computational power to perform these tasks (Figure 2).

Machine learning (ML) is a subfield of AI that studies algorithms and techniques for automating solutions to complex problems that are hard to program with traditional programming methods. In traditional AI programming, a computational program is given a task without instructing the machine on how to solve it. The first step in this method is for the programmer to develop a detailed plan for the program to solve the target program, to define a set of rules applicable to solve the problem. Then the specific detailed plan is programmed in a computer language (Rebala et al., 2019). However, developing a plan to solve many real-life problems is not always manageable. For example, let us say we want a machine to detect different English accents in audio inputs. A dataset of audio is provided, and each audio data is labelled with a specific English accent – New Zealand, Canadian, and so on. The labelled datasets are the set of examples that the machine will train on for the computational program to recognise a new accent provided, such as Indian English. In a conventional program, the programmer needs to study all the audio inputs to understand how different accents correspond to English speech varieties and to devise a detailed plan consisting of a set of rules for the machine to detect these accents. This task can be a very challenging task to complete.

ML algorithms are designed to solve real-life problems. In other words, they are designed to learn from the dataset. Therefore, the larger the dataset, the better they learn. Therefore, the ML algorithm indirectly solves a problem by first generating a model based on processing the dataset and then predicting the label of new input data by executing that model. The output decision relies on a two-fold process: on the input data and the algorithms selected by the experimenter to detect relationships; and the machine's automatic learning ability. Machines need data to learn, to perform tasks efficiently and to generate outcomes quickly (Badillo et al., 2020). The uniqueness of ML consists in the fact that the model and the human operate apart, in contrast to traditional programming (Theobald, 2017).

Notably, the definition of learning in ML differs from that used in behaviour analysis. At the same time, learning in behaviour analysis is the underlying process of behavioural change, contingent upon changes in reinforcer manipulanda over time. As ML is an AI system, it should be able to learn in a changing environment, for example, handling more data (Alpaydin, 2020). If the machine learns to adapt to environmental changes, it should address all sorts of problems, by acquiring the ability to make decisions based on the data and the experience gained from previous learning (Alpaydin, 2020). However, unlike behaviour analysis, where the goal is to understand the underlying mechanisms of learning, in ML, the goal is to develop a functional, well-engineered system.

ML algorithms run on the paradigm of neural network systems, called connectionist systems, composed of complex mathematical computations with multiple single processes (denoted as units), connected to each other to resemble the human brain's neurons. An *artificial neural network* (ANN) is a brain-inspired network model that simulates the brain in analysing and processing information. Deep learning is a specialised field of learning that uses advanced algorithms simulating multilayered neural networks to process large datasets.

The third-generation ANN, known as *spiking neural network* (SNN), employ spiking neurons as computational units. Spiking neurons are mathematical models or differential equations, that describe much better the actual output of a biological neuron than older ANNs did. In addition, they are energy-efficient like neurons that generate spikes of membrane action potential of a tiny voltage.

More on SNN in relation to other models will be described further in this chapter (2.5).

Figure 2

A visual overview of artificial intelligence (AI) and its subfields

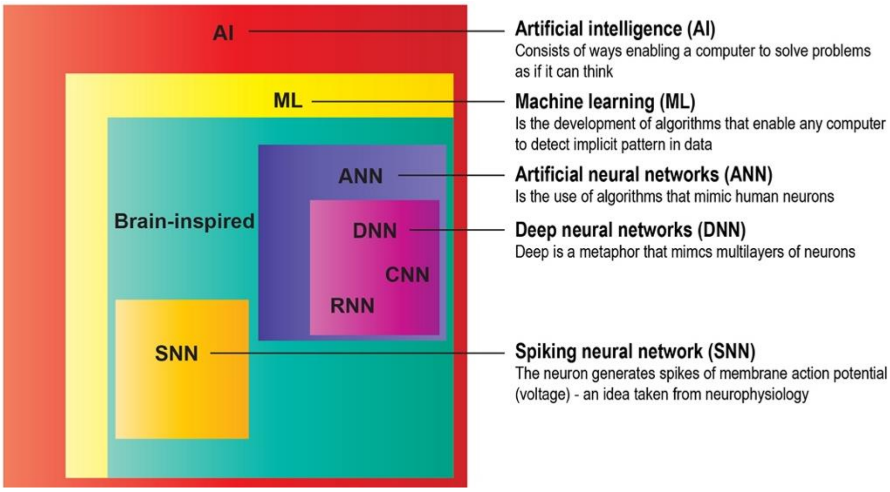

2.2 ML and Psychology

As introduced above, ML is a type of AI constructed to improve itself continuously through processing data. It addresses the issue of solving problems by learning from data, particularly exemplars. ML algorithms allow software applications to make accurate predictions. Successful applications indicate that these computational models can independently learn, grow, develop, and adapt.

Advancements in ML have allowed scientists to experiment with ML methods across scientific domains, commerce, and technology, leading to alternative pathways for even more evidence-based predictions (Badillo et al., 2020; Jordan & Mitchell, 2015, Kuhn & Johnson, 2018). In the area of healthcare, four main applications of ML have been reported in the literature:

(i) detection and diagnosis;

(ii) prognosis, treatment and support;

(iii) public health; and

(iv) research and clinical administration.

The most common mental health conditions addressed include depression, schizophrenia, and Alzheimer's disease (Shatte et al., 2019). Emphasis has been given to creating and improving existing ML models for disease detection to minimise the burden on health (Verma & Verma, 2022). Another area of enormous interest has been the early detection and diagnosis of individuals with autism spectrum disorder (ASD). However, despite models reaching accurate predictions in detecting ASD, these models have shown to have difficulties in real-life applications, which may be related to the data selected to train the models. This possibly has caused misalignment with the observational basis upon which professionals diagnose ASD (Cavus et al., 2021; Kassraian-Fard et al., 2016). Besides diagnosing ASD, today we find several applications of ML in mental health (e.g., Asif et al., 2020; Justice et al., 2019; Malgaroli & Schultebraucks, 2020; Ratana et al., 2019; Vaidya et al., 2020).

Other applications, for example, are in education, by creating intelligent tutorial systems (Korkmaz & Correia, 2019). They are also useful in animal research, to address evolutionary questions about animal behaviour (Valletta et al., 2017); with electrophysiological data for predictions (e.g., Karekla et al., 2022), and demonstrating many real-world applications (Sarker, 2021). In an attempt to understand the human brain and, by extension, human behaviour, recent papers have recommended that the next phase in research should be simulation neuroscience, conducted through the construction of models of the brain (Fan & Markram, 2019).

Dwyer et al. (2018) have proposed a theoretical framework for assessing ML generalisability across individuals and clinical psychological settings, to reinforce the potential of ML applications. Moreover, Yarkoni and Westfall (2017) argued that the progress in ML technologies has made it feasible to generate efficient translational results in clinical practice. Lastly, it has been suggested that ML analysis may assist psychology experiments in developing more efficient techniques for analysing results and minimising replicability issues (Orrù et al., 2020). Some recent studies in psychiatry and psychology have emphasised that we need to be cautious about how machine-produced predictions can be interpreted or even translated into clinical practice (Meehan et al., 2022; Volovici et al., 2022). Clinical data present with unique characteristics (e.g., heterogeneity) and false positive outcomes can significantly impact our therapy design. For ML to be an accurate and useful predictive tool in the psychology toolbox, cautious action must be taken. Most research has focused on improving ML performance. Though ML performance needs to be assessed, the focus must also shift onto assessing ML usage to assist clinical decision-making (Ehrmann et al., 2023). The latter has been researched even less.

This thesis originates from an assumption made from the very beginning that evaluating ML performance is insufficient. For example, most problems focus on improving its accuracy through experience that consists of input-output pairs (Jordan & Mitchell, 2015). Therefore, a computational model's outcome may be evaluated positively for making accurate predictions in relation to the whole dataset. However, it may fail when the outcome is evaluated by how well it performed for all positive predictions that it could possibly make. Following the same logic, taking all sorts of performance measures would not solve the problem. Therefore, careful steps need to be taken to assess the performance that suits the data and the problem to solve. An

ML model's utility increases when the predictive ability gained, complements the mathematical models built on 'control'.

2.2.1 Contribution of ML to Behaviour Analysis (EAB and ABA)

ML can arguably contribute to laying the foundations for new, applied training methodologies using optimal training regimes for specific learning goals and specific individuals. These training methodologies can help us better understand adaptive learning mechanisms in frequently changing environments. Furthermore, they can contribute to creating predictive models of learning depending on environmental parameters. Experimental research could then be translated more easily and provide better, personalised interventions for those in need. In addition, research has recommended that 'good modelling practices' should encompass the evaluation of a model, its generalisation ability, conditions for comparing data, and model comparability (Cox, 2019).

EAB has a new, unique chance to advance its technologies with the use of AI algorithms.

2.3 Machine learning for behaviour analysis

A few ML applications have been utilised within the EAB field. For example, machine learning has been used to interpret behavioural phenomena such as *autoshaping* and *automaintenance* (positive and negative) (Burgos, 2007). The study was designed to simulate the two phenomena, with neural networks limited to two features. Burgos (2003) also applied machine learning to simulate latent inhibition. Both studies involved multiple simulations to investigate different conditions and neural-network architectures, by integrating operant and respondent conditions. Overall, the studies demonstrated the use of neural-network models to understand

and simulate complex phenomena in behaviour analysis, including autoshaping, automaintenance, and latent inhibition. Furthermore, Burgos developed three different four-layer feedforward neural networks that did not differentiate between operant and Pavlovian conditioning while retained the standard of behavioural distinctions. The author demonstrated the application of a neural network model in studying *misbehaving* by integrating the control of respondent and operant conditioning in specific moments (Burgos, 2015), thereby illustrating the relationships between different types of conditioning and behaviour.

The overarching goal in the above-mentioned studies was to illustrate how computational modelling can learn more about biological constraints to learning and formulate how knowledge can be gained about the former by examining the latter, without one contradicting the other. Furthermore, Burgos' research aimed to test a computational modelling of conditioning in situations where performances cannot be tested with actual participants, to enrich our understanding of the distinction between operant and respondent learning and the relationship between behaviour and neuroscience (Burgos, 2010) rather than to predict future behaviour (Burgos, 2018; Ninness et al., 2018, 2019; Ninness & Ninness, 2020). The authors were interested in the predictive ability of neural networks and they investigated how, by using simulated *matching-to-sample procedures*, an ANN learns these relations like humans would do. They further demonstrated how machine learning algorithms predicted the relationship between two stimuli that were not taught directly but rather were inferred by a matching-to-sample procedure (i.e., *derived relations*). Moreover, simple artificial models have been tested for the study of *stimulus equivalence* to study symbolic behaviour, discussing that artificial models can enhance the prediction of human behaviour through simulation procedures (Tovar et al., 2023). Finally, Tovar et al. discussed how computational models can simulate learning equivalence relations.

Conversely, the learning process via *operant conditioning* has inspired computer science researchers to use ANN to act as robot brain controllers (e.g., Cyr et al., 2014). Based on the principles of adaptive learning in behaviour analysis, Cyr et al. investigated whether a machine learning model (SNN) could make a robot change its response contingent upon changes in the environment depicted by changes in the delivery of consequences. They demonstrated that an adaptive learning process within a bio-inspired paradigm (SNN) was possible with a simple neural architecture for achieving complex behaviours. This and the above-mentioned studies demonstrate that machine learning has the potential to be applied to and learn from behaviour-analytic experiments, which opens a bidirectional path between EAB and computer science.

Limited studies have been published that focus on applying machine learning to the clinical practice of learning theory (i.e., ABA). ML has been tested in measuring challenging behaviours. For example, a recent study applied neural networks to measure the duration of vocal stereotypy, as observing such a behaviour is resource-intensive (Dufour et al., 2020). Dufour et al. concluded that the ML model was able to identify with high-accuracy vocal stereotypes, such as repetition of sounds, words, or phrases, from audio data with and without background music, for six out of eight participants diagnosed with ASD. Another study applied ML to detect and distinguish among different types of self-injurious behaviours, such as head hitting, self-biting, eye gouging. They used body-worn sensors for individual and group data and received high scores of ML performance across different measures (Cantin-Garside et al., 2020). Overall, research in this area is still primitive.

Additionally, other ML applications aided behaviour assessments in identifying the function of a problem behaviour (Functional Behaviour Assessments). For example, one study applied ML to find ways to assist with a behaviour assessment – either to improve the interpretation of visual graphs (Lanovaz et al., 2020), or to identify the function of behaviour when using a questionnaire about the function of the problem behaviour (QAFB) (Bailey et al., 2021). These studies indicated that the machine outperformed any other methods used. In another study, hypothetical data compared the ML model's performance to human performance, in the case of special education teachers. The teachers viewed simulated intervention graphs and selected a clinical decision based on the data. The study supported the idea that humans' decisions were significantly improved with an algorithm, providing a focused summary of data, and thereby showing possible ways in which ML could be helpful in clinical decision-making (Brodhead & Truckenmiller, 2021). Another area of interest was using ML to identify and predict treatment outcomes in an intensive early intervention programme for young children with ASD (Linstead et al., 2015, 2017). Linstead demonstrated that the ML model could predict mastery of learning criteria outcomes and that higher treatment intensity predicted significant progress.

Lastly, ML algorithms have also been used to extract information on the scientific evidence in published behavioural research; for example, by identifying trends in the behaviour-analytic literature (Sosine & Cox, 2023) or identifying misconceptions about behaviour analysis in non-English posts on the internet (Turgeon & Lanovaz, 2021). A summary of the rising trends is presented in Figure 3.

Figure 3

A visual summary of topics of interest for the use of ML in behaviour analysis

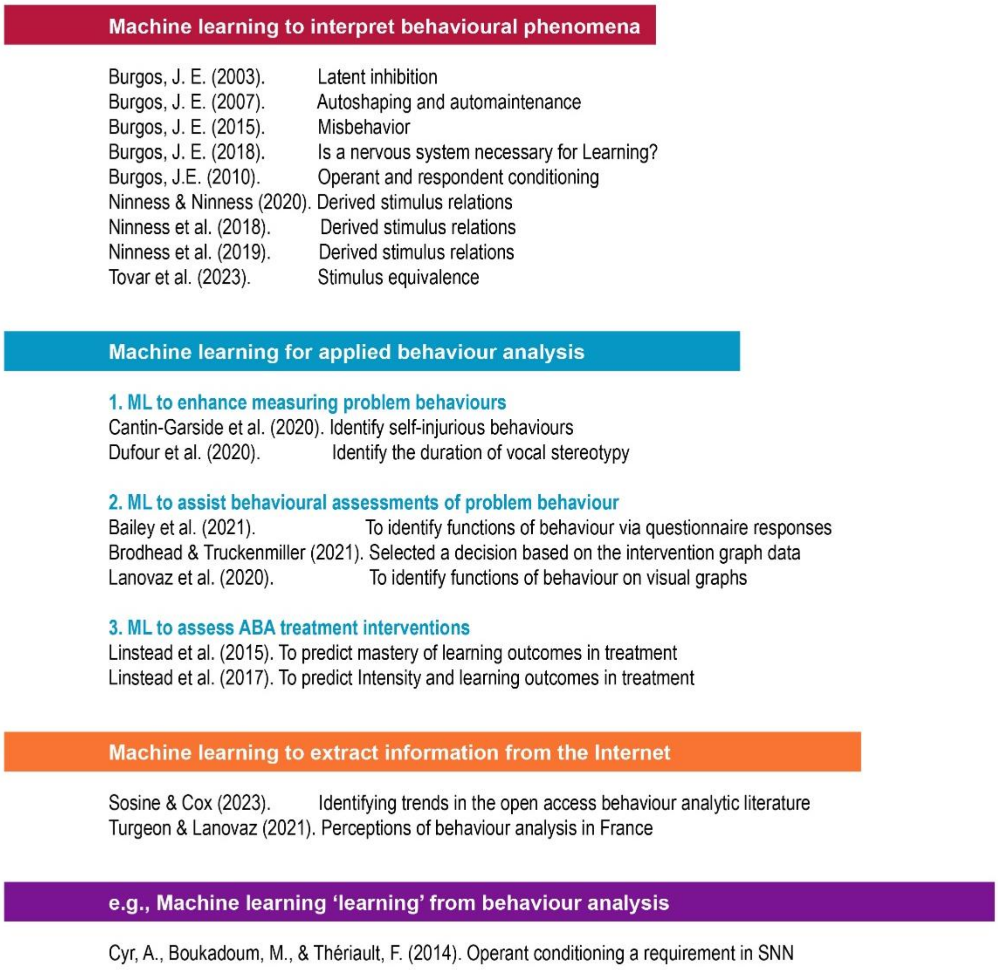

Machine learning to interpret behavioural phenomena

Burgos, J. E. (2003). Latent inhibition
Burgos, J. E. (2007). Autoshaping and automaintenance
Burgos, J. E. (2015). Misbehavior
Burgos, J. E. (2018). Is a nervous system necessary for Learning?
Burgos, J.E. (2010). Operant and respondent conditioning
Ninness & Ninness (2020). Derived stimulus relations
Ninness et al. (2018). Derived stimulus relations
Ninness et al. (2019). Derived stimulus relations
Tovar et al. (2023). Stimulus equivalence

Machine learning for applied behaviour analysis

1. ML to enhance measuring problem behaviours
Cantin-Garside et al. (2020). Identify self-injurious behaviours
Dufour et al. (2020). Identify the duration of vocal stereotypy

2. ML to assist behavioural assessments of problem behaviour
Bailey et al. (2021). To identify functions of behaviour via questionnaire responses
Brodhead & Truckenmiller (2021). Selected a decision based on the intervention graph data
Lanovaz et al. (2020). To identify functions of behaviour on visual graphs

3. ML to assess ABA treatment interventions
Linstead et al. (2015). To predict mastery of learning outcomes in treatment
Linstead et al. (2017). To predict Intensity and learning outcomes in treatment

Machine learning to extract information from the Internet

Sosine & Cox (2023). Identifying trends in the open access behaviour analytic literature
Turgeon & Lanovaz (2021). Perceptions of behaviour analysis in France

e.g., Machine learning 'learning' from behaviour analysis

Cyr, A., Boukadoum, M., & Thériault, F. (2014). Operant conditioning a requirement in SNN

2.4 Conclusion: ML can advance behaviour analysis and vice-versa.

The studies discussed above constitute early examples of how AI tools can be added to the behaviour-analytic toolbox. Furthermore, their utility was demonstrated for both research (EAB) and clinical practice (ABA), opening avenues for predictions of all sorts.

An advantage of using an ML model is that it allows the experimenter to distance themselves from the data, as the model itself detects implicit patterns in the data. For example, from the mathematical models used in choice research, it has been observed that behaviour proportional to the relative reinforcer rate could sometimes produce different predictions based on the different mathematical models used (Mazur, 2006). This generated different assumptions about the learning process that underlined the observed behaviour. Therefore, an artificial model using algorithms might assist the researcher in making a decision when comparing outcomes of different mathematical models. In other words, AI can assist the researcher in gaining knowledge from data, in the process of making informed decisions.

A second advantage of using ML in behaviour analysis is for EAB research to create opportunities to disseminate knowledge to the broader society. This can be achieved by providing those who offer services to people in need, such as clinical behaviour analysts, ABA therapists, behavioural psychologists, and teachers, with access to the sophisticated knowledge accumulated by EAB research. Computer scientists, neuroscientists and biobehavioural researchers would equally benefit from such access. Discussions have been held and it has been pointed out that despite rigorous experimental research results, EAB is not creating a social impact, and there is still a gap between EAB goals and real life (Killeen, 2018). For example, an applied behaviour analyst, may find it essential to detect an observed behaviour's learning history, as it is often not available in natural settings and is crucial information for treatment design. The information is necessary to design a learning environment for the behaviour to change positively, from maladaptive to adaptive. EAB researchers may not be motivated to answer a question such as what is this past learning history, as they are the ones that fit more with the traditional experimental approach and manipulate environmental changes. There is subjectivity to what someone may find interesting to study, and a selection bias in the choice of research questions that will influence what will be studied (Critchfield, 2011). In addition, a practising behaviour analyst’s enjoyment of mathematical equations can vary. Understanding published EAB research can often require a relatively sophisticated understanding of both psychology and mathematics. In contrast, AI tools can extract the necessary knowledge for applied usage without the user having expert knowledge, thus allowing EAB research to become friendlier and usable for specific purposes. Therefore, AI can assist EAB research in addressing an applied problem without abandoning its focus on fundamental principles.

Furthermore, animal experiment design and analysis can become time-consuming and require much effort to translate outcomes into clinical practice. AI tools may assist in creating collaborations and knowledge-transfer pathways. It has been argued that designing animal experiments with automated measurements, is no longer the norm within the overarching behavioural science (Locey, 2020). Some behaviour-analytic researchers are choosing other means of doing science, some by using indirect methods (e.g., Burney et al., 2023) or correlational designs (e.g., Young, 2018). These deviate from the fundamental approach in traditional methods, that is, inductive, within-subject approach, for direct measures to be taken and replicated systematically. Though this translates in the science reaching out for applications to build social impacts, there is also a greater risk regarding the fidelity of the procedure of replication of results (Locey, 2020). In this thesis, the argument is not to divert from the traditional approach to conducting research but to demonstrate to the researcher or the practitioner that AI tools can complement, assist, or have utility in the analysis of results. EAB datasets can fulfil this function with a twofold purpose: to increase social impact while carrying on expert research in developing learning theories.

Finally, with the continuous advancement of AI technologies, the future may have ML models learning with nonhuman animals, like pigeons. Thereby, AI tools will allow the design of animal experiments to take place without involving animals at all, where the artificial models are implicitly learning from data and experience. For example, a recent study trained an ML model to learn the participants' choice-making strategies by imitating the process underlying the subjects' choices and their learning abilities (Dezfouli et al., 2019). In robotics, a team built an operant conditioning paradigm where the robot made real-time choices (Cyr et al., 2014) or switched behaviour in response to a threatening event (Cyr et al., 2020). In summary, EAB knowledge of adaptive learning can be used to potentially create pathways for collaborations to advance adaptive learning methods and theory in ways that were once only part of science fiction.

AI may assist the field in keeping behavioural science a purely natural science by continuing to generate experiments that control events in the environment and by asking new research questions beyond the initial research designs by analysing carefully produced datasets with the help of AI. As every dataset presented to an AI tool, among other things, affects the machine's ability to learn, this thesis has used behavioural datasets extracted from EAB choice experiments. The behavioural paradigms developed for choice research are scholarly and sophisticated, and patterns in behaviour reflect explicit environment arrangements.

SNNs –third-generation artificial neural networks– have demonstrated the ability to capture spatial and temporal dimensions in the data (more in Section 2.5). Furthermore, contemporary research in choice behaviour has acknowledged that the spatial and temporal properties of reinforcer-response relations can assist in understanding learning mechanisms (Cowie & Davison, 2016a). Hence, in this project, SNNs were tested on EAB choice behaviour datasets. The following section explains why SNN was the ML of choice to be tested with these datasets.

2.5 Spiking Neural Networks

2.5.1 Introduction

As presented in the Introduction of this chapter, an Artificial Neural Network (ANN) is a brain-inspired network model that simulates the brain to analyse and process information. Processes for artificial neural networks (ANN) have been studied for more than fifty years and are the building blocks of deep learning architectures. The human and nonhuman brain inspires these ANN which are designed to mimic their behaviour with the ambition functioning as and understanding a living organism's brain. ANN are connectionist systems with complex mathematical computations. Computer scientists have created artificial-like neurons connected to each other in a net to simulate the human brain (van Gerven & Bohte, 2017).

Unlike human learning, ANN learns from datasets in three different ways. First, supervised learning would require the data (inputs) to come with labels (for example, reinforcer condition type), and the task would be for the machine to be fed with a label input. Learning occurs when an error signal between the actual data and the label is calculated by adjusting the network weights to minimise the error and, thus, adjusting the parameters. Second, when learning is unsupervised, the data fed come with no labels; with no feedback being provided about the models' performance. Third, when both learning procedures are combined, semi-supervised learning occurs. Lastly, in reinforcement learning, the model is the decision-making agent that acts in an environment that receives a reward or a penalty for its actions to solve a problem (Alpaydin, 2020).

ANNs are classified into three generations according to their computational units. The first generation utilises the McCulloch-Pitts neuron (also known as *perceptron* algorithm), a mathematical model where information passed is represented as numerical values. They are universal models for computing digital input and output. Some examples of this generation are the *multilayer perceptron*, *Hopfield nets*, and *Boltzmann machines* (Maass, 1997). In the second generation, we can find mathematical models that apply an *activation function,* a mathematical computation, added at the end of the network to determine what information should pass on to the output. When the mathematical model applies the activation function, it uses a continuous set of possible output values to a weighted sum of the inputs, which are the real values that correspond to important input data characteristics. Unlike the first generation that produced digital inputs and outputs, the mathematical model can now compute *analogue* inputs and outputs, continuous signals (Alaloul & Qureshi, 2020). Another characteristic of the second-generation ANNs is that they support learning algorithms based on *gradient descent*, an algorithm that optimises the result, such as *backpropagation*. Typical examples of these networks are *feedforward* and *recurrent sigmoidal neural nets* and networks of *radial basis function* units.

The recent success of AI is attributed to deep learning, a branch of ML (Figure 2). Deep learning has gained wide adoption on the internet by companies such as Apple, Google, Microsoft, and more (e.g., Alzubaidi et al., 2021; Khurana et al., 2023) . The main successes in psychology research lie in examples such as brain image analysis, emotion recognition, speech audio, and natural language processing, including the analysis of social media posts (e.g., Lu, 2022) . Despite the popularity of DNN architectures, deep learning is not a single approach but rather a class of algorithms and topologies that can be implemented to make decisions in response to a range of questions. According to the literature, each architecture built could have its own strengths and weaknesses, and there are no clear one-size-fits-all solutions (Alzubaidi et al., 2021)

The third-generation, also called spiking neural networks (SNNs), employ spiking neurons as computational units (Figure 4), which are mathematical models that describe much better the actual output of a biological neuron (Gerstner & Kistler, 2002; Maass, 1997). One advantage from previous generations is that the algorithm is inherently suitable for temporal data and has a vital role in the computational method (Izhikevich, 2004). Thus, spiking neurons have an *all*-or-*none* response where all information is encoded in the spikes' timing, whilst the neuron's dynamic nature changes with time (Ghosh-Dastidar & Adeli, 2009). Therefore, unlike the two previous generations, which trivialised the timing of computation steps, the timing of individual computation steps plays a crucial role in computations in networks of spiking neurons. This is achieved by using *spikes timing codes*-discrete events placed at a point in time-as the main carrier of information rather than continuous changing time values (Rafi, 2021). Therefore, SNNs are computationally more powerful than conventional continuous or rate-based ANNs as they are time-defined (Maass & Markram, 2004; Tavanaei et al., 2019).

Figure 4

A schematic representation of spikes as the main carrier of information for the spiking neuron

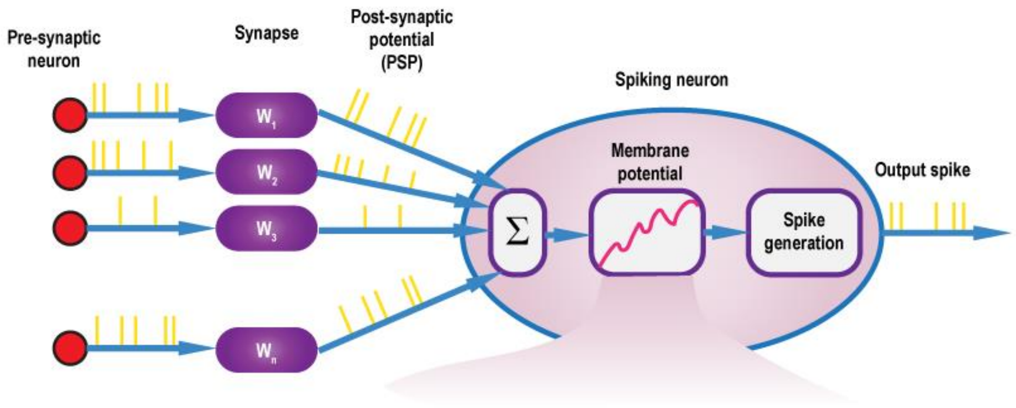

The occurrence of a spike is determined by *differential equations* that represent various "biological" processes of the machine, the most important of which is the neuron's membrane potential. Essentially, once a neuron reaches a certain potential, it spikes, and the potential of that neuron is reset. Additionally, SNNs are often sparsely connected and take advantage of specialised network topologies (Soni, 2018). Therefore, SNNs learn from spatiotemporal data and are capable of making generalisations to new data (Kasabov, 2019). However, despite their ability to solve complex problems by leveraging the timing of discrete spikes, SNNs have difficulty because there is no single way to encode the information into spike events. As a result, the signal process that carries the information for the neuron to learn can vary. This has slowed the process of SNNs finding real-life applications (Auge et al., 2021).

SNNs offer an additional advantage in that they do not require *batch training*, unlike ANNs. This means that SNNs can adjust their weights much faster than ANNs, as they do not need to train the model on subsamples of data at a time, as is done with the gradient descent algorithm commonly used in ANNs (Cyr et al., 2014). For example, an SNN, when trained with unsupervised learning, can learn from the initial pair of spikes in just one trial. Hence, they demonstrate in this example a critical ability to learn in a real-time context which cannot afford learning on slow synaptic adjustments as in rate-based models (Cyr et al., 2014). Therefore, SNNs seem capable of handling complex dynamic information, including binary choices, within a given time frame.

2.5.2 Computational model of a spiking neuron

Biological neurons are the excitable cells of the nervous system. They connect to each other through synapses forming a neuron network that encodes for the processing of information in the brain. Neural signals propagate via action potentials (APs), meaning that through electrical activity, a stimulus will cause the neuron to send information down the axon and away from the soma. APs are electric impulses or changes in the membrane potential that travel along the neuron surface. When a threshold potential is reached, then AP is generated (voltage increases); neuroscientists call this response a spike. Neurons connect to each other via synapses, which are specialised structures enabling the formation of functional networks that process and store information (for more, (Krebs et al., 2018). For example, a human brain can classify the name of an object, such as a cup, by identifying salient features of the object regardless of size, colour, shape, or material made. How the brain works and manages such tasks has inspired computer scientists to move away from traditional computer paradigms and to mimic (or at least be inspired by) biology to create several neural models.

A computational spiking neuron produces spikes that are discrete events that take place at points in time, coming from one neuron to another and connected through weights that act as synapses (Maass, 1997). This is achieved with complex mathematical computations using different spiking encoding techniques to model spikes. Spiking neurons carry and transmit information through the generation of electrical spikes (or action potentials). Inspired by how information is transferred in the biological neuron they consider the spike's exact time or the spike or the sequence of the spike (Figure 5). As presented in the example, the model learns by using biologically plausible algorithms (Σ), feeding the information forward, as in a natural brain, across the neural network, starting from the input node to the output node. The real-world data are decomposed into input vectors. The weights represent the synaptic connections of neurons and refer to the strength of those connections, i.e., the effect of the firing rates. An input current stimulated the neuron model to produce a sequence of spikes with a specific firing rate. Based on weights, the input signal may be amplified or inhibited. When the membrane potential reaches a certain threshold, the neuron will spike, generating a signal that will travel and instantly reset to a lower value. The neuron stays in a resting state for some time, called the absolute refractory period, after which it can process new information coming from other presynaptic neurons (Gerstner & Kistler, 2002). Thus, this process resembles the natural nervous system and is considered biologically realistic and plausible (Maass, 1997). In essence, SNNs' strength is that they have an event-driven asynchronous operation by using spikes, the networks learn from spatiotemporal data and, critically, they have the ability to re-learn and adapt the connections with other neurons while already working, akin to how the brain works (Khadeer, 2020).

Besides the rationale that SNN could provide a possible good fit with small-behavioural datasets, I conducted preliminary performance tests of other ML technologies by using the pigeon datasets used to train the SNN model in this study. The data were classified by three machine learning techniques: Multilayer perception (MLP) – a feedforward artificial neural network (LeCun et al., 1998); Support vector machine (SVM) – a supervised learning model that has been used extensively for classification that works by finding a hyperplane in a high-dimensional feature space that best separates different classes of data points (Cortes et al., 1995); and Evolving classification function (ECF) and EC with Clustering the data (ECMC) – ECF is based on an evolving connectionist system learning from input data and evolving its connections by extracting meaningful rules that are applied on new input data (Kasabov, 2019), thus adapting itself to changing conditions by updating with new data and ensuring accuracy in non-stationary data distributions; ECMC takes this adaptability further by integrating evolving classification with clustering techniques, enhancing adaptability to diverse patterns within the dataset. For all three traditional ML methods, cross-classification was conducted using k-folds, as with SNN. A comparison of performance was conducted between the traditional machine learning techniques and the performance of SNN models built in a NeuCube architecture using a NeuCube java application on Netbeans IDE 8.2, that includes several algorithms to build a SNN model. The NeuCube architecture uses a brain-inspired environment of SNN for data mining, pattern recognition and predictive data modelling with spatio-temporal data (Kasabov, 2014) and was available for use by researchers at AUT KEDRI. All traditional machine learning techniques (MLP, SVM, ECF, ECMF) and NeuCube were trained and tested on the same pigeon datasets as described in Chapter 4. The comparative results of these classification tests are presented in Table 1.

Table 1

Overall accuracy of cross-validation classification results from the comparative analysis of four traditional ML models and four SNN models

	SNN Algorithms				Traditional ML models			
No of Samples	**TBR$_{thr}$**	**AER**	**BSA**	**SF**	**MLP**	**SVM**	**ECF**	**ECMC**
175 samples (5 pigeon datasets)	0.72	0.86	0.65	0.84	0.52	0.32	0. 54	0. 57
210 samples (6 pigeon datasets)					0.28	0.41	0. 59	0.57

Note: The SNN models were trained on 175 samples; further generalisation tests were conducted with the remaining 35 samples.

Results from traditional machine learning techniques were generated from one dataset that was split for training and testing each developed model. Firstly, the cross-validation was executed with the 175 samples that were used with NeuCube for cross-validation training/testing procedure. Secondly, the same procedure was executed with all 210 samples, as the classification results in traditional machine learning varied in relation to the sample size; more samples produced a somewhat better accuracy.

Additional generalisation results were not obtained as the 35 samples used in NeuCube were too small to generate an outcome with the traditional ML models. The change in the number of input samples affected positively only the SVM and ECF results, yet it did not catch up to the accuracy of SNNs. In an attempt to improve MLP accuracy, a further 10 experiments were executed with both sets of input data; the average results for training and testing improved, however not sufficiently so, as accuracy scores varied between 0.57–0.59. The architecture of the java version of NeuCube displayed the best performance in terms of accuracy when compared to other computational models. Given that these results clearly appeared to be more advantageous than the results of the traditional models and that the goal of this doctoral thesis was to capture the dynamic and event-based nature of spatiotemporal data while incorporating biological realism, the SNN appeared to have more value for the pursuit of my research. I decided to attempt engineering an SNN model of as simple a structure as possible to test its potential for accurate prediction based on our datasets (see Chapter 4).

Figure 5

The visual shows a representation of a biological (top) and brain-inspired computational neuron (bottom).

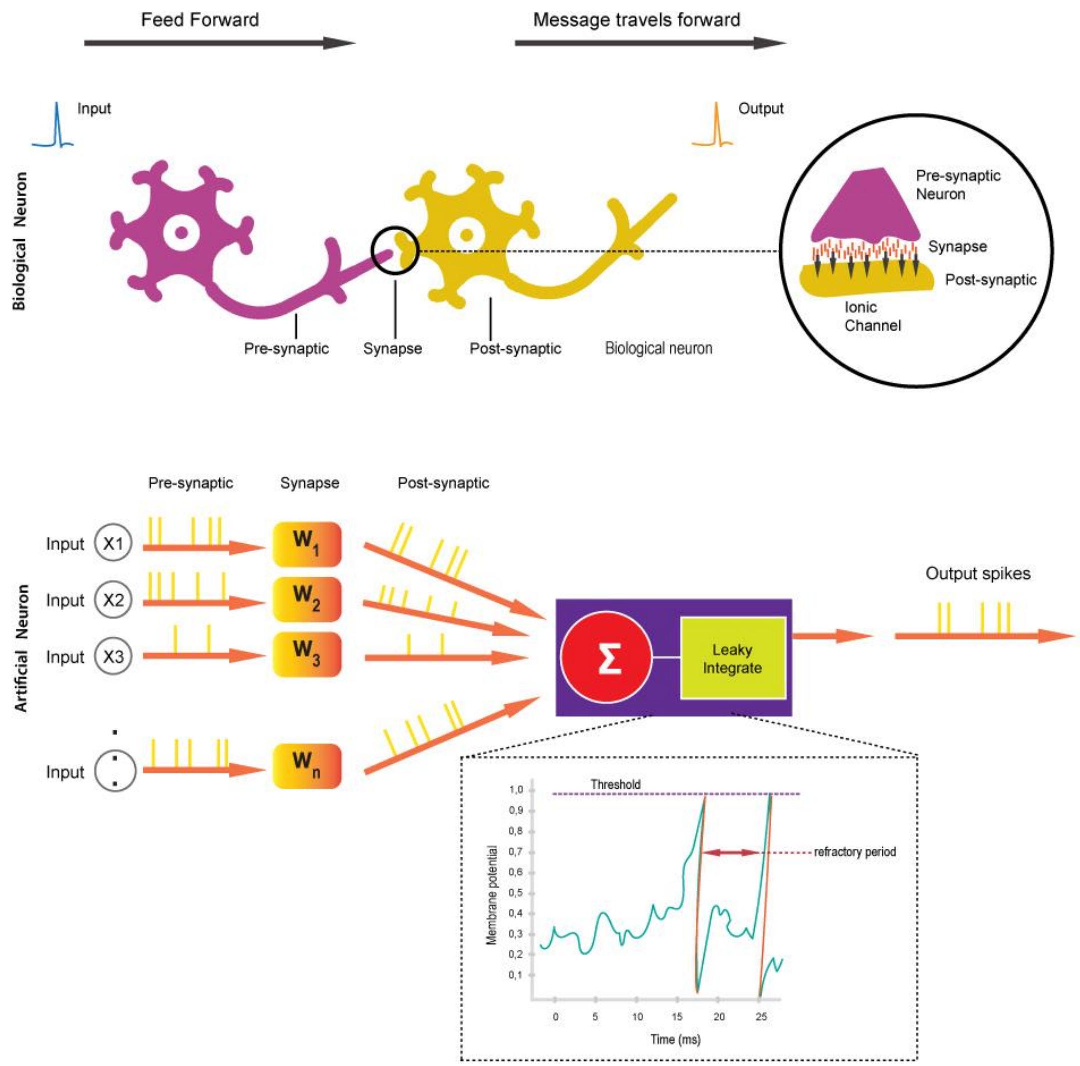

A spiking neuron fires whenever its membrane electric potential reaches a certain threshold. The membrane electric potential is the sum of excitatory and inhibitory postsynaptic potentials, which result from the firings of other neurons connected through the weights. The firing of a presynaptic neuron at times contributes to the membrane's electric potential, and the neuron's mathematical model describes the amount.

Artwork is adapted here and was designed for the experiment in Chapter 7.

Leaky integrated-and-fire (LIF) spiking neuron model

The Leaky Integrate-and-Fire (LIF) is the most common computational neuron model that defines the evolution of the membrane potential but also the mechanism to generate spikes with low computational cost (Gerstner et al., 2014). Specifically, a LIF is a mathematical abstraction that takes input currents into the neuron at different times of postsynaptic potentials that have a subsequent effect on the membrane potential (Gerstner & Kistler, 2002). Thus, it lays its foundation on how the membrane integrates the excitatory post-dynamic currents until the membrane potential reaches a threshold for spiking (Burkitt, 2006). It is most widely used due to its simplicity and accuracy in modelling biological neural network dynamics (Tan et al., 2020).

The LIF neuron used for this project was developed in (Gerstner & Kistler, 2002) and denoted as the *leaky* LIF as it is modelled as a "leaky integrator" of its input I(t), which simulates the sum of excitatory and inhibitory postsynaptic potentials; The LIF behaviour is modelled using a resistor-capacitor (RC) circuit where the leakage is due to the resistor and the integration of I(t) is due to the capacitor that is in parallel to the resistor, adhering to the following equations:

$$\frac{d}{dt}(\upsilon) = \frac{[RI(t) - (\upsilon t - \upsilon rest)]}{\tau m} \qquad (1)$$

where υ denotes the membrane potential, R the membrane resistance, I(t) for instantaneous currents, $\upsilon rest$ for resting voltage, and τm is the time membrane constant calculated as:

$$\tau m = RC \qquad (2)$$

The spiking events are not explicitly modelled in the LIF model. When the membrane potential υt reaches a certain threshold υth (spiking threshold), a spike is generated as:

$$s(t): \upsilon t \geq \upsilon th \qquad (3)$$

The equations are a simplified version of the organisation and neurophysiological processes of the biological neuron. In particular, the spatial structure and the mapping of presynaptic to postsynaptic potentials are put in the weights. However, it is still a functional simulation (Cyr et al., 2020). To add some realism to the dynamics of the LIF model developed for this thesis, an absolute refractory period Δabs immediately after υt hits υth was added. During the absolute refractory period, υt might be clamped to $\upsilon rest$ and the leaky integration process is re-initiated following a delay of Δabs after the spike (for more, see section 4.1 (*4.1 Integrate-and-Fire Model*, n.d.)).

In conclusion, due to the deterministic nature of the leaky LIF model, it is the model of choice to train the model to identify the learning histories from the datasets. As mentioned above, the occurrence of a spike is determined by differential equations and a differential evolution (DE) algorithm for this thesis. The DE algorithm (DE/rand/1/bin) is a population-based heuristic method for global optimisation over continuous spaces (Storn & Price, 1997). The Leaky LIF was trained to identify the learning histories by adapting the weights w and generating the input current I(t) that characterises each learning history. This approach is addressed as an optimisation process, where the objective is to find the optimal w values that maximise the prediction accuracy. The model was considered because it is a robust and straightforward algorithm, as its search process is governed by a few algorithm-specific parameters (Ahmad et al., 2022). The DE algorithm has been proven to be effective in solving various problems while showing flexibility and stability as a population-based algorithm. Research has shown that despite DE simplicity, even the standard DE manages to outperform other adaptive schemes regarding the quality of the result and the convergence speed (Georgioudakis & Plevris, 2020).

2.6 Behaviour datasets for Artificial Neural Networks

In ML identifying patterns would require data and, therefore, datasets. A dataset comprises multiple data points of different natures, and each point represents an entity for analysis (Badillo et al., 2020). Therefore, data points can be numerical, categorical, time-series or discrete events and extracted from various sources. The quality and quantity of the data play a significant role in the training of the algorithm. Poor quality datasets, for example, wrong labels or inconsistent data, can lead to incorrect predictions. Furthermore, data can be biased, leading to biased predictions (Tato & Nkambou, 2022). Also, *Imbalance*[1] datasets can make predictions difficult (Johnson & Khoshgoftaar, 2019).

The 'Big Data' era has attracted many, including behaviour analysts, to focus on large datasets (Cox et al., 2021). Of course, the 'attraction' comes with a reason, as deep neural networks need a considerable amount of data to be trained. The essence of looking into small datasets to capture the heterogeneity has also been discussed in psychiatry in respect of ML model use (Koppe et al., 2020). Also, the data quality is easier to monitor than in large quantities, which, as stated above, can affect the model's predictive ability. ML has shown that a number of approaches can be applied to small datasets even when the number of variables used for such a task outnumbers the participants (Orrù et al., 2020). The rationale for using small datasets is explained below.

---

[1] Imbalance dataset would imply that there is a disproportion of data points for different examples, thus making the number of each set of examples different (e.g., more datapoints related to the function of attention vs escape in a behavioural assessment of challenging behaviour)

Lastly, there is a growing interest in finding applications of ML models with non-imagery or non-physiological datasets, as most models developed today rely on imagery or physiological data with limited applications to other forms of data (Dehghan et al., 2022).

2.6.1 Spiking neurons to model small-N behavioural datasets

Psychological research mainly implements two methodologies: group and small-number (small-N) designs. Most fields of psychology rely on comparisons between groups, with only a few exceptions, such as EAB and one or two others (Hayes, 1981). As discussed in Chapter 1, EAB is committed to analysing strictly observed behaviour of individual organisms, though more than one individual is studied. Consequently, these two methodologies differ in their approach: group designs use between-group comparisons and rely heavily on inferential statistics, whereas small-N designs mainly use within-subject designs and rely strictly on observable phenomena, with a built-in emphasis on replication and reproducibility with other subjects (Sidman, 1960). Traditionally, group designs are used in cognitive psychology and small sample designs in behavioural psychology. However, EAB research is considered to generate replicable outcomes by the very nature of its designs (Imam, 2018). Most papers supporting ML for psychological research favour group designs, different from this study, which is based on smaller-N designs constituting an original research approach.

In this thesis, we are using small-N design data for several reasons. Firstly, ML applications are at an early stage in psychology, and it is vital to investigate the quality of the constructed model. As different approaches to ML may lead to different outcomes, all ML outcomes need human interpretation, which can only be achieved by evaluating the model itself (Cearns et al., 2019). Furthermore, our model results can be compared to the current choice research results. Therefore, we need to rely on soundly researched datasets that do not depend on p-values (Ioannidis, 2016, 2018). Based on Cearns et al. (2019), if we can extract critical variables from an ML model while retaining performance, we may achieve theoretical understandings equal to or even better than those obtained from methods currently employed with small-N designs. At the same time, it saves resources on all levels.

Furthermore, as dependent variables within small-N designs are constantly manipulated and repeatedly measured in each phase of an experiment (Kazdin, 2011), temporal and spatial data forming a causal relationship between variables, can be used meaningfully. Every choice behaviour is a datum, and the advantage of the neural network of choice in this study -SNN- is that it can process events without the need to average them. In the process of an individual organism learning, changes occur; orderly change is only seen when we average the data from many trials or many organisms, but with SNN, we do not need to present the data in an orderly fashion. Thus, the importance of a single organism is not lost, and more meaning can be extracted from the data (Skinner, 1984b). According to Skinner (1984), it is not enough to record the onset or endpoint of a behaviour since we can miss meaningful information from the continuous process. For example, choice research has shown that previous reinforcer deliveries affect subsequent behaviour, and the dynamics of this relationship are determined by environmental variation and the frequency of environmental changes influence the speed at which changes in choice behaviours occur (Davison & Baum, 2000). Therefore, the effects of individual reinforcers differ in a number of ways depending on the context in which they are obtained; for example, the reinforcer rate arrangement is the alternative that produced the previous reinforcer and the number of successive reinforcers obtained from alternatives

(Landon & Davison, 2001). If a series of cross-sectional data is taken from a specific point of time, important information about the degree of environmental variability can be missed. In this example, how fast or slow would learning have been, depending on the reinforcer's delivery? Moreover, depending on the number of environmental variables, how would short- and long-term learning have been affected? Small datasets are used as tasks and experimental methods need to discriminate between conditions that are still under investigation.

When creating a model, it is important to be able to answer questions like: "How can we be sure that the artificial model models choice behaviour?", "What are the relevant features of the environment that can be captured in the SNN model?", "How accurate can the predictions be, and how well can they be generalised to other data?" and, in more general terms, "How is the SNN handling these variables and what patterns does it extract?"

SNN is a new subdivision of neural networks that offers spatiotemporal data processing. As the elapsed time and the location of choice change (left to right and vice versa), what role do these parameters have in choice patterns? In this thesis, learning is conceptualised, based on functional relations between learning histories and changes in choice behaviour. SNN technology can extract data from temporality and discover patterns and produce spikes that take place at points in time. The analysis of spatiotemporal variables is fundamental to this conceptualisation. It allows the examination of data variability and whether patterns persist over time or if any unusual patterns emerge over a particular period. Small-N designs are ideal for capturing these purposes for studying individual organisms and the significance of small datasets has been raised in other domains as well (psychiatry; (Koppe et al., 2020), neuroimaging; (Vabalas et al., 2019); clinical prediction; (Riley et al., 2020)).

2.7 Overall summary for Chapter 1 and Chapter 2

In Chapter 1, we appreciate that quantifying observable behaviour, is a bottom-up approach for psychologists to learn from the data rather than by making theoretical assumptions. The mathematical models adopted by those devoted to viewing the analysis of behaviour as a natural science, made it possible to measure behavioural changes in their interplay with the environment. This led to the development of an ongoing understanding of adaptive learning and to efforts to predict single organisms' behaviour based on observed instances. However, prediction based on reinforcer deliveries has been proven to be complex, and research findings show that simple learning processes may be more complicated than initially presumed. We cannot just accept that if reinforcers are given, the same behaviour will recur under similar environmental conditions as implicated by the Law of Effect (Ferster & Skinner, 1957; Thorndike, 1911).

Research on choice behaviour, has demonstrated that by strictly controlling the manipulation of environmental variables (e.g., rate of reinforcement, the magnitude of reinforcer, time-elapsed, sequence or reinforcers,), we observe multiple factors influencing whether choice behaviour will or will not be demonstrated. This control operates on a multilevel (e.g., localised vs long-term) (Landon & Davison, 2001). Contemporary research has demonstrated that the past, present and future are important elements in prediction (e.g., Cowie, 2018; Cowie & Davison, 2020a). When an organism shows preference by making a choice, that choice behaviour is not solely based on response-reinforcer associations. Instead, the organism acts like a computational agent, and every time it needs to make a choice, it computes from the input data — the variables affecting behaviour in the past— to decide what the future outcome will likely look like (Cowie & Davison, 2020b).

It has been argued that if we move away from the fact that the reinforcer just induces behaviour (Baum, 2012), we may better understand these environment-behaviour interactions and, thus, predict behaviour. However, today, new predictions are not attainable if the learning history is not available. The field of artificial intelligence may be the place to look for alternative tools to detect the learning history when this history is unknown to us. Temporal variables such as choice behaviours may allow us to derive conclusions about learning that took place in a specific environment.

The discussion in this chapter demonstrated new ways to conduct research and determinants that influence the interpretation of the outcome. Psychology discussion papers suggest that artificial intelligence models can be successfully used to understand organisms' behaviour. To date, models as such have successfully been used in education and the health sector and have shown clinical benefits when machine learning (ML) methodologies are considered. The benefit of using ML models, is that it allows us to step away from existing relations, such as response-reinforcer and learn from the data to make predictions without the need for previous knowledge. Its experimental nature relies on identifying unique patterns residing in the input data using complex computational techniques. The bottom-up approach adopted in artificial neural networks seems to fit well with EAB. Thus, any previous knowledge that is derived from the relation between input, including choice behaviour and class labels., such as reinforcer arrangement, is not necessary. The relationship is identified solely on pattern detection and prediction tasks rather than previous knowledge on choice research. The advantage of making novel predictions would imply that the unknown information on the learning history can be identified. Thus, ML opens new avenues to ask new questions on prediction without restrictions.

Furthermore, the output decision of an ML model does not solely rely on the input data but instead on a twofold process – the input data and the algorithm selected, which enables the machine to learn. Thus, it is necessary to test if an ML model can work with small behavioural datasets and determine in what modelling context predictions can happen. Research has demonstrated ways in which performance for accuracy in predictions can be tested, yet most clinical research applies to image data (Dehghan et al., 2022). If behavioural research datasets are to be used, we need to take for granted that the datasets are small and diverse, as behavioural patterns are shaped by the environment and by individual differences, which are determined by biological factors. Thus, how behavioural data can be used as inputs and the type of neural networks would need careful attention when building machine learning models.

Over the years, SNNs have shown interesting developments. Several advantages of using SNN have been described in this chapter. However, a crucial advantage of SNNs over other ANN is that the neuron learns from spatiotemporal data incorporating the timing of spikes in the algorithm and have the potential to make generalisations to new data (Yamazaki et al., 2022). Choice behaviours are discrete events of observable data that have been shown to have a complex relationship with reinforcers exhibiting spatial and temporal properties. SNN technology was chosen as it was hypothesised that event-based processing might be more suitable for choice behaviour datasets because it allows outputs to be derived from spatiotemporal information and discover patterns without aggregating or averaging the data. This allows us to examine variability in the data and see if patterns persist over time or if any unusual patterns emerge over a particular period. Secondly, it allows the 'Big Data' approach to be avoided because data from fewer individuals may be more important to a model's accuracy when applied to a single person.

2.8 Conclusions in a nutshell

SNN might allow us to get a unique insight into past experiences. The quantitative nature of the data from EAB choice experiments provides reliable datasets necessary for the machine to learn and we may well get some solid outcomes from there.

Most research on ML today focuses on the brain to understand behaviour. However, behaviour analysts are interested in understanding behaviour from an environmental perspective – how the environment brings about changes in behaviour.

This doctoral thesis aspired to explore a method of analysis of behavioural spatiotemporal data (from EAB research) to identify learning histories. The results can contribute to laying the foundations for new applied training methodologies using optimal reinforcer arrangements for specific learning goals and specific individuals. Furthermore, if modern technology allows us to challenge our teaching and training methods by changing the way we ask questions and analyse problems, we may be able to understand learning as a situation-specific problem-solving strategy (Skinner, 1984a) and develop learning markers that predict future learning patterns. Experimental research could then be translated easier and provide better, personalised interventions for those in need of them.

Chapter 3 Rationale and significance

3.1 This thesis

Behaviour is understood in relation to past experience; therefore, current patterns in behaviour should allow prediction of past experiences. If an SNN model is able to predict the latter (training environment) from the former (current behaviour), it would mean that the behaviour carries information about the environment, confirming the existence of their inter-relation and generating a new hypothesis for studying environment-behaviour interactions. Such predictive outcome would be of value to clinical/ABA treatments and interesting from a theoretical perspective.

In clinical practice, the learning history of a target behaviour is commonly unknown and can only be speculated upon to inform assessment and treatment; so most often, little attention is paid to the matter. Having the ability to access learning histories would better inform us in what conditions a new behaviour can be learnt. For example, a child may use their past experience to find ways to disrupt the classroom in order to obtain immediate consequences (e.g., peer attention). Therefore, knowing the patterns of learning (i.e., the reinforcer contingencies previously in effect) might point to effective training arrangements that can produce the desired future behaviour.

From a theoretical perspective, an SNN may allow us to get a unique insight into past experiences, should it be capable of ‘predicting’ learning histories. The predictive outcome can be investigated further in the context of other stimuli taking place in the environment. For example, differences in predictive outcomes may help us understand behaviours that have been controlled by various aspects of the environment. For example, behaviour may inform us about the temporal structure of the environment (such as a more remote learning history contributing more (or less) to current behaviour patterns).

The quantitative nature of choice behaviour datasets used in this thesis made them an excellent candidate to be transformed into spikes. Pigeons' choice behaviours were made in environmental conditions that changed rapidly. They were the best candidate datapoints as the environment contained elements of variability that often feature in natural environments (in contrast to steady-state choice procedures), and the research had demonstrated that choice behaviours adapted to their environment quickly (Davison & Baum, 2000; Landon & Davison, 2001). It was important to use a dataset that resembles the natural environment as it is a challenge to identify the unknown learning histories (Lattal, 1995). At the same time, only a fraction of the experimental data was included in the experimental application of this new technology, which constitutes another novel factor introduced in this thesis research project. In particular, the datasets constituted of binary choices that occurred for a 5-sec period after the delivery of each reinforcer. The patterns of responses immediately after the reinforcer are often referred to as *preference pulses* (Cowie & Davison, 2016a). Local analyses of choice behaviour that occur after a reinforcer delivery has shown that reinforcers are followed by a brief change in the response ratio towards the just-reinforced alternative (e.g., Davison & Baum, 2002; Landon et al., 2002, 2003b) . However, a participant may choose in certain environments to peck the same key rather than to switch to another key for reasons other than the immediate reinforcer delivery (i.e., short-term effects) and the pattern of its responses to be controlled by the long-term effects of reinforcers (McLean et al., 2014). Thus, regardless of the analysis of how patterns of choice behaviour relate to reinforcement (i.e., either associated to the immediate reinforcers or with some event in the distant future) the existence of the relationship should make it possible for a ML model to infer the specific learning history based on current behaviour. Furthermore, it is crucial to consider data taken from short time periods (i.e., 5-sec), as very often in clinical situations (i.e., contexts outside laboratories) it is not practicable to collect large amounts of data under strictly controlled conditions. This is crucial to consider when attempting to translate research into clinical practice.

An artificial spiking neural network (SNN) was chosen, due to its ability to process discrete events that occur in time by using spikes rather than having to treat the data as a cumulative number of responses over a continuous period. Thus, the advantage of the algorithm is that it inherently suitable for temporal data. SNNs also allow for the spatiotemporal analysis of the data considering the behaviour allocation (left-key versus right-key) from one time to the other regardless of how the relationship between behaviour and reinforcer was formed and does so by developing a nonlinear model.

This thesis has the potential to demonstrate how SNNs (spiking neural networks) can fit with behaviour datasets, such as choice behaviour data, and learn patterns from observed behaviour under various reinforcer arrangements. This offers a new tool for research methodologies. Lastly, in this thesis, we wanted to strengthen the EAB’ s translational goal by using datasets to answer questions (e.g., what is the learning history of the choice I am observing) that have both a clinical/applied and basic research value and may bridge the gap between the two fields –experimental and applied– in behaviour analysis.

3.2 Overall contribution of this thesis

Throughout this thesis, these items have been discussed:

- This thesis demonstrated that it is viable to approach data differently by using a snapshot of data and relying on observable behaviour and relevant processes to make predictions rather than manipulating the reinforcer-behaviour relation, thus circumventing lengthy common training procedures.
- This thesis added a tool in the behaviour analysis toolbox that can predict the unknown learning histories from limited datasets containing a small 'window' of binary choices.
- This thesis demonstrated that the artificial model was adaptive and responsive to individual organisms' data. Hence, predictions can be personalised and used to study individual learning patterns, which is the goal of behavioural research.
- The prediction of learning histories lays the foundations for a novel approach to examine the use of new training methodologies, using optimal training conditions for specific learning and for specific individual organisms.

3.3 Rationale, objectives, research questions

The development of AI in recent decades has allowed us to use this technology and advance our ways of doing science. This thesis aims to advance behaviour analysis by demonstrating how an AI tool can empower the researcher to ask questions that could not be easily considered with the existing tools in the EAB toolbox and without specialised knowledge in choice research experiments. The research objective was not to find the best artificial model by tweaking engineering structures and techniques or developing algorithms. Instead, it was to find out whether an SNN is useful with behavioural datasets and whether we can have a decision by the machine (i.e., a prediction) that makes sense, is reasonable and usable.

The two main questions underpinning this thesis are:

1. First, if only a snapshot of choice behaviour was available, can an SNN predict the learning history (in other words, the reinforcer arrangement) used in training conditions that led to the current choice behaviour known to us?

2. If yes, what happens if we want to predict the learning history of other participants or learning histories that differ in the reinforcer delivery from the ones the AI model used for its development?

Table 2 at the end of this Chapter presents the research key points and links between experimental studies in the following chapters.

3.4 Thesis structure

This thesis combines a published paper, two articles submitted for peer review, and a study ready for submission, to address the research questions of this thesis. According to each journal's specifications, each article was written for psychologists or a computer science audience.

The novelty of the approach and the limited research in the area required some background for the reader to understand why the prediction of the learning histories is an important question and how AI technology can assist in answering such alternative questions. Therefore, an understanding of the objectives and research questions of this thesis were provided in Chapters 1 and 2. Chapters 4 and 5 focus on answering the first research question, and Chapters 6 and 7 focus on the second research question. A prelude section was added to each of Chapters 4—7 to provide critical complementary information, that could not be included in the submitted papers to assist the understanding of a psychologist or computer scientist respectively, when the article is not directly related to their field. In addition, linking paragraphs between Chapters 4—7 were added to make the overarching narrative of this thesis more seamless. As this thesis required the researcher to develop new skills to make this research possible, a short list of skills and knowledge gained is also provided at the end of Chapter 4. Finally, the papers in Chapters 4, 5 and 7 were initially prepared according to each journal’s author guidelines. For consistency and clarity in this thesis, English spelling, table and figure numbering, and referencing have been adjusted. Table 2 presents critical research points from the experiments developed for this thesis and the links between them.

Finally, Chapter 8 presents a general discussion that summarises the overall findings of this thesis, the contribution to behavioural analysis, limitations, and future directions. In the end, a personal reflection on the PhD process is provided. Figure 6 presents a schematic flow of the thesis.

At the outset of this thesis, the plan was to link Chapters 1 and 7 with a study that would combine the outcomes from rat and human datasets. The aim was to assess how variability in the data itself may affect the machine's ability to predict their learning history when predictions made for one species are made from an artificial neural network trained with data from another species (e.g., pigeon datasets). Some pilot experimentation was conducted after the study in Chapter 6 (Appendix C). Unfortunately, this study was not completed as planned because of COVID-19, time delays and employee layoffs that led to changes in the specialised assistance I received for the SNN development from an AUT research fellow (Dr Israel Espinosa-Ramos). Due to the uncertainty around the amount of the work needed and the timeframes that this would require, together with the supervisors we decided that the body of work completed should be the subject of this thesis, whereas the inter-species comparative study should be put on hold for future research. Nevertheless, it will be discussed in the future directions section in Chapter 8.

Figure 6

A schematic flow of the thesis.

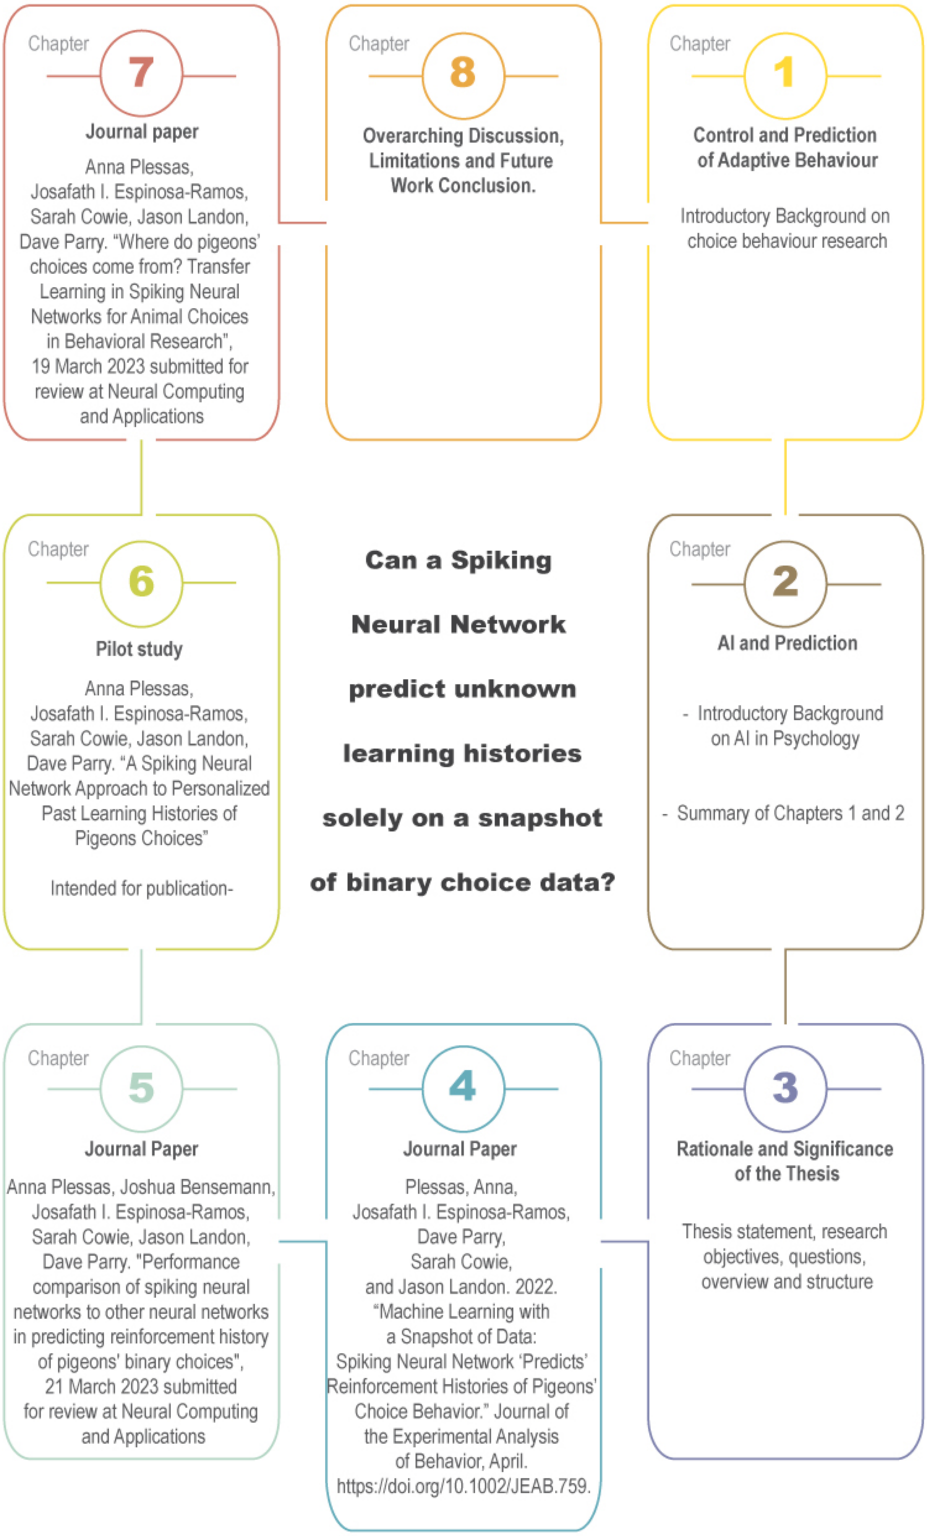

Table 2

Research Key Points and Links Between the Chapters of Experimental Studies

Overall Question 1: If only a snapshot of choice behaviour was available, can an SNN predict the learning history-the reinforcer arrangement used in training conditions- that led to the current choice behaviour known to us?

Chapter 4

Modelling learning history: The application of a spiking neural network model on pigeon's choice behaviour

Chapter Content

Question:	**Rationale:**	**Novelty in the Approach:**
When the organism's learning history is unknown, can an SNN model containing one LIF neuron identify this history based on the current choice (behaviour)? a) When snapshots of binary choices are available, and b) The dataset of snap-shots of choices is small	An accumulated body of choice research has demonstrated that choice can be understood within the context of the learning history of each choice. Researchers typically arrange different 'histories' by varying reinforcer conditions and observing the resultant behaviour to understand how learning history affects behaviour. By the same logic, it should also be possible to identify the specific learning history based on the current choice when this history is unknown.	Experimental Study 1) A single-neuron SNN was trained to learn patterns from current binary choices data coupled with specific learning histories (i.e., concurrent variable-interval (VI) schedules). 2) Tested whether this model could predict specific learning histories with small datasets containing a snapshot of binary choices (5-s periods after the delivery of the reinforcer).
Findings:	**Main Novel Contributions to Behaviour Analysis:**	**Secondary Novel Contributions to ANN:**
• Pigeons' binary choices provided sufficient information for the single-neuron SNN to predict the specific learning history with a good degree of success without including other data from the learning history of the animal. • The SNN outcome matched the analysis of the pigeons' actual choice, supporting the idea that SNN can process behavioural data to make predictions when the learning history is unknown. • The SNN generalised to novel datasets of the same participants to identify the specific learning history when the total number and size of datasets varied. Thus, the ability of the SNN model to make predictions was related solely to the pigeons' unique response pattern and the SNN's ability to learn.	• The study demonstrates that a novel ML tool (i.e., a single-neuron SNN) can inform behaviour-environment contingencies using a snapshot of data without using standard testing-training procedures. • Also, the study demonstrates that SNN can solve specific problems by circumventing lengthy common training procedures. • Moreover, the study demonstrates how the predictive outcome can be utilised as a hypothesis generator directing us to look into specific data for further analysis and investigation. • Lastly, this study lays the foundations of a novel approach to examine new training methodologies, using optimal training conditions for specific learning and specific individual organisms. The prediction outcome can help the behaviour analyst's decision-making process when designing such environments to change behaviour and shape future choices by creating learner-specific and task-specific environments.	• The study demonstrates a temporal input data stream with an SNN model using one LIF neuron for the first time.

Links between Chapter 4 and Chapter 5

Having developed an SNN model that could conduct a reverse test to predict learning histories solely on binary choice datasets, validating its effectiveness by comparing it with other types of ANNs was essential. Different neural network architectures can affect the predictive outcome. Thus, before experimenting further with the developed SNN, using the same datasets on other ANN was important to see if we had created a good enough model.

Chapter 5

Performance comparison of spiking neural networks to other neural networks in predicting the reinforcer history of pigeons' binary choices.

Chapter Content

Question:	Rationale:	Approach:
Can deep learning models produce predictive outcomes as the developed SNN model by Plessas et al. (2022) did to identify the unknown learning history of pigeons based on their current choice behaviour? a) Can RNNs (LSTM, GRU) and CNNs produce a predic- tive outcome? b) How is their performance when compared to the SNN developed in Chapter 4?	In Chapter 4, the single-neuron SNN demonstrated that it is possible to conduct a reverse test to predict the unknown past training environment by relying solely on known observed behaviour (choices). The approach was novel in that a reverse test was conducted using SNN and small datasets of snapshots of binary choices shaped across time. Event-based processing used in SNNs was suitable for choice data processing because it allows outputs to be derived from spatiotemporal information rather than from averaging the data. This study assessed the performance of the same datasets with deep learning ANNs that have been demonstrated to handle sequential data. Performances were compared to validate the predictive outcome.	Experimental Study

Findings:	Main Novel Contributions to Behaviour Analysis:	Secondary Novel Contributions to ANN:
• All models (SNN, GRU, LTSM, CNN) could predict the pigeons' learning history with an overall average accuracy of 89%. • When other measures were considered, such as F1, the SNN outperformed the prediction ability of all other DNNs. The result indicated that all DNNs would benefit from some optimisation while the SNN did not require further training and could successfully handle the data.	The study demonstrates that: • Event-based processing, used with SNN, fits with observable binary choice datasets well. • Temporality is a crucial parameter with binary choices shaped in frequently-changing environments. • SNN is the easiest option of new-tech tool additions to the toolbox of behaviour analysis.	The study demonstrates: • A novel investigation of reverse tests with different types of neural networks with real datasets. • The approach can help to guide the development of more effective neural networks.

Overall Question 2: What happens if we want to predict the learning history for other participants or learning histories that differ in the reinforcer delivery from the ones the AI model used for its development?

Links between Chapter 5 and Chapter 6

Different neural network architectures differ in predictive ability. After demonstrating in Chapter 5 that the single-neuron SNN developed in Chapter 4 could reliably be used with choice research datasets, the next step required further investigation with other datasets. A useful SNN model should be able to be applied to other participants' choices (either groups or individuals) doing the same task, even when choices were learnt in similar but not identical environments. By including datasets from other studies, external validation of the SNN model application is strengthened, and the predictive outcome allows us to investigate further environment-behaviour interactions.

Chapter 6

A spiking neural network approach to personalise learning histories for pigeons' choice-making

Chapter Content

Question:	Rationale:	Approach:
When generating predictions, can the single-neuron SNN model generalise to predict the unknown learning history of other participants that display new individual differences in choice patterns? a) When using datasets from other participants, whose learning was shaped through learning histories that are identical or non-identical to the ones the single-neuron SNN was trained on b) Can small datasets constituting a snapshot of binary choices produce a personalised predictive outcome? c) How is the predictive performance affected by group datasets and individual datasets?	Learning is a unique process for each organism. Any method or tool developed for studying data from behaviour-environment relations cannot be a one-size-fits-all technology. In the present study, the previous research findings were extended to investigate whether a single-neuron SNN model can be personalised by considering the unique patterns developed by each organism during learning.	Experimental Study
Findings:	**Main Novel Contributions to Behaviour Analysis:**	**Secondary Novel Contributions to ANN:**
• When the outputs of the single-neuron SNN (i.e., firing rates) were redefined with the new datasets, the predictive outcome was higher than traditional approaches to generalisation. • In particular, when using new choice data from other participants, the model responded to the new dataset and predicted the learning history both for each individual and for the group. However, overlapping patterns of choices with reinforcer conditions that were closer to each other made it harder to detect a unique pattern. • Moreover, when using new choice data coupled with non-identical environments —variation in the number of the reinforcers given and the speed of change in reinforcers given— the predictive outcome remained high. This outcome indicated that the differences in learning histories did not affect the model's ability to discriminate patterns.	The study demonstrates that: • The single-neuron SNN predictions can be tailored specifically to individual learning differences that represent data variability. • The outcome is trustworthy, as the SNN model was adaptable and data-responsive. • The SNN can be a valuable tool for studying single organisms' behaviour, as the SNN outcome aligns with EAB research. In addition, the outcome can support the behaviour analyst's decision-making when designing environments to fit specific individuals or groups to change their future choices.	• The study demonstrated a transfer learning technique with an SNN model to overcome data dependency, generalise new datasets and personalise the output.

Links between Chapter 6 and Chapter 7

The study in Chapter 6 was a pilot study, and the method used was a first step towards experimenting with generalisation issues when using new datasets. The study in Chapter 7 focused on assessing generalisation again by refining the method used in the pilot study. A range of transfer learning methods was used to evaluate how transfer learning can be applied without having to train the SNN model from scratch every time a new dataset is used to predict the learning history of an organism based on the organism's binary choices. It allows us to investigate further generalisation issues with small datasets.

Chapter 7

Where do pigeons' choices come from? Transfer learning in spiking neural networks for animal choices in behavioural research

Chapter Content

Question:	Rationale:	Approach:
Can a single-neuron SNN model generalise to novel instances of the same or similar problem? a) When nested cross-validation is used b) Is prediction better with any of the two transfer learning training strategies (feature extraction and fine-tuning)?	For an SNN model to find applicability in behaviour analysis, it needs to have the ability to deal with small datasets and to generalise to new datasets that represent variability due to individual differences in learning. Nested cross-validation is a training technique capable of estimating a model's generalisation error and the hyperparameters search. The estimation of errors is achieved by using a series of train/validation/test set splits into the data, creating two cross-validation steps. By refining the SNN training method and comparing more than one transfer learning strategy, the SNN performance was assessed against all sorts of binary choice data to predict learning histories. Therefore, all datasets used were the same as in Chapter 6.	Experimental Study

Findings:	Main Novel Contributions to Behaviour Analysis:	Secondary Novel Contributions to ANN:
• All SNN models developed with nested cross-validation could handle different datasets (mean score 83%). As in Plessas et al. (2022), the predictive outcome was higher for binary choice shaped in extreme reinforcer condi-tions (i.e., majority of reinforcers delivered on one key) than for those shaped in nearby reinforcer conditions (i.e., reinforcer delivery looks similar). This indicates that similar patterns in choice can exist in two similar learning histories but the further these histories are, the more distinct the patterns in choices. Yet, the SNN model's predictive performance varies depending on the dataset's characteristics used for training or testing. • When comparing both transfer learning strategies, fine-tuning slightly outperformed feature extraction. However, this does not indicate that one strategy is better than the other, as both produced comparable predictive outcomes. • Feature extraction allows prediction to be produced within seconds in contrast to fine-tun- ing, which needed more time as retraining parts of the model was required. Predictions after fine-tuning required half the time that retraining the model with nested cross-validation requires.	The study demonstrates that: • Both transfer learning strategies are suitable for studying animal behaviour, with outcomes consistent across participants and experimental procedures. • As transfer learning methods can reach similar or better predictions than traditional approaches to generalisation, behavioural analysis can use a tool that allows fast results and predictions on low-budget devices. • Pattern detection is feasible for various small datasets consisting of binary choices. • The SNN model is adaptive in targeting the data of individual organisms, thus allowing the study of individual organisms, which is the goal of behavioural research.	• The study demonstrates a transfer learning approach with non-imaging and non-physiological data. • Contribution to the sparse research in transfer learning with a spiking neural network.

This research was unfolding during the COVID-19 pandemic. Thus, scheduled participation to present the research at International Conferences was cancelled. Instead, the research was presented to the University of Auckland Neurocognitive lab, KEDRI and AUT students co-authoring papers on AI.

3.5 Research publication resulting from this doctoral thesis

Plessas, A., Espinosa-Ramos, J. I., Parry, D., Cowie, S., & Landon, J. (2022). Machine learning with a snapshot of data: Spiking neural network 'predicts' reinforcement histories of pigeons' choice behaviour. Journal of the Experimental Analysis of Behaviour. https://doi.org/10.1002/JEAB.759.

3.5.1 Research publication currently under review

Anna Plessas, Joshua Bensemann, Josafath I. Espinosa-Ramos, Sarah Cowie, Jason Landon, Dave Parry. "Performance comparison of spiking neural networks to other neural networks in predicting reinforcement history of pigeons' binary choices", 21 March 2023, submitted for review at Neural Computing and Applications

Anna Plessas, Josafath I. Espinosa-Ramos, Sarah Cowie, Jason Landon, Dave Parry. "Where do pigeons' choices come from? Transfer learning in spiking neural networks for animal choices in behavioural research", 19 March 2023, submitted for review at Neural Computing and Applications

3.5.2 Other publications in preparation

Plessas A, Espinosa-Ramos J, Landon J, Cowie S, Parry D. A spiking neural network approach to personalise learning histories for pigeons’ choice-making. To be submitted in June 2023

Plessas, A, Weerasinghe, M, Parry D. The application of a spiking neural network with temporal neurobehavioural data. To be submitted

Plessas A., Espinosa-Ramos J., Landon J., Cowie S, Parry D. Can a Spiking Neural Network inform us on Past Learning Histories Across Species? To be submitted

Chapter 4 Modelling learning history: The application of a spiking neural network model on pigeons' choice behaviour

Prelude

As discussed in Chapter 1, experimental analysis of behaviour (EAB) has extensively studied choice behaviour, and in recent times a line of research has focused on choice in variable conditions. When shaping choices in rapidly changing procedures, reinforcer conditions change multiple times in a session in an unpredictable manner. This procedure simulates the natural environment to some degree and results in behaviour that changes quickly in response to the changing contingencies. EAB researchers observe the resultant behaviour to understand how reinforcement history affects behaviour. These experiments occur in a laboratory, often using caged pigeons as participants, and research has demonstrated that animal choices conform to the same general rules as human choice behaviour (e.g., Cowie et al., 2021).

At the commencement of this Chapter, there is a section about how datasets from choice research experiments were used in this project to train a simple and efficient, brain-inspired AI tool, namely an SNN containing just one LIF neuron. This section provides a description of how choice experiments are conducted, what constitutes a reinforcement history in such experiments, and lastly, the procedure used to train the pigeons and the outcomes of the source experiment (Landon & Davison, 2001). This information was not included in the published paper presented in this Chapter and is included to provide a broader context for the work.

Choice behaviour and experimental arrangements

In choice research, behaviour analysts pay attention to how the experimenter arranges the training environment. When conducting choice experiments, the environment consists of *reinforcers* usually arranged using schedules of reinforcement. There is constant interaction between organism and environment whereby environmental changes are followed by the organism responding to them, thereby changing the environment, and so on. This enables every procedure to have a feedback function that can be described mathematically or graphically to demonstrate how the environment behaves as a function of an organism's behaviour (Baum, 1973). Therefore, in order to understand learnt behaviour, it is important to control and systematically observe how changes occur during this dynamic interaction.

A common procedure used is a *variable interval schedule of reinforcement* (VI) presented using a *two-key concurrent schedule procedure.* Variable interval schedules provide reinforcement for a target behaviour (key pecking here) when a variable time has elapsed since the last reinforcer (Ferster & Skinner, 1957). In a VI, the variable time is set and has a fixed mean and distribution,

$$R = \frac{1}{t + 0.5/(\frac{1}{B})} \qquad (1)$$

where R is the reinforcer rate, B is the response rate, and t is the scheduled time value.

The functional relationships between the environment (R) and the behaviour (B) on a VI schedule would imply that the reinforcer rate between two reinforcer deliveries, equals the scheduled time interval plus half the average time responses distributed between two reinforcer deliveries (1/B) (for more, (Baum, 1973). It has been studied that the performance of a VI schedule can be affected by several factors, such as the level of deprivation the animal has, the reinforcer rate or magnitude. This, in turn, can affect the rate of the animal's responses. Furthermore, in concurrent VI schedules, responding is shaped by two or more separate VI schedules simultaneously manipulated by the experimenter. For example, a concurrent VI 27-s vs VI-27-s means that two keys for responding are available, and a VI 27-s is arranged for each key (Figure 7). The VI runs at the same time for both keys, even if the animal is only responding to one key until the reinforcer is delivered and the time stops until the reinforcer has been taken. VI can be arranged in two ways. One way is to arrange them independently of one another (e.g., Herrnstein, 1961). Another arrangement is to stop the timing for both Vis when a reinforcer is delivered. When the reinforcer is collected, timing is resumed for both VI (Stubbs & Pliskoff, 1969). The pigeon chooses between the two alternatives in a two-key concurrent VI schedule in a way that the relative frequency of responses on one key equals the relative frequency of reinforcer delivery for the same key (Herrnstein, 1961). *Changeover delays* are necessary with concurrent VI schedules and used in choice research to reduce the possibility that the gain from the reinforcer obtained on one alternative, is passed on to the other key. When using a changeover delay parameter in a concurrent VI schedule, the reinforcer cannot be obtained until a short period (2-3 seconds) has elapsed since a changeover, even if it has been arranged by the schedule (Shull & Pliskoff, 1967).

Figure 7

An example of a two-key concurrent VI schedule manipulation for a choice research experiment with caged pigeons

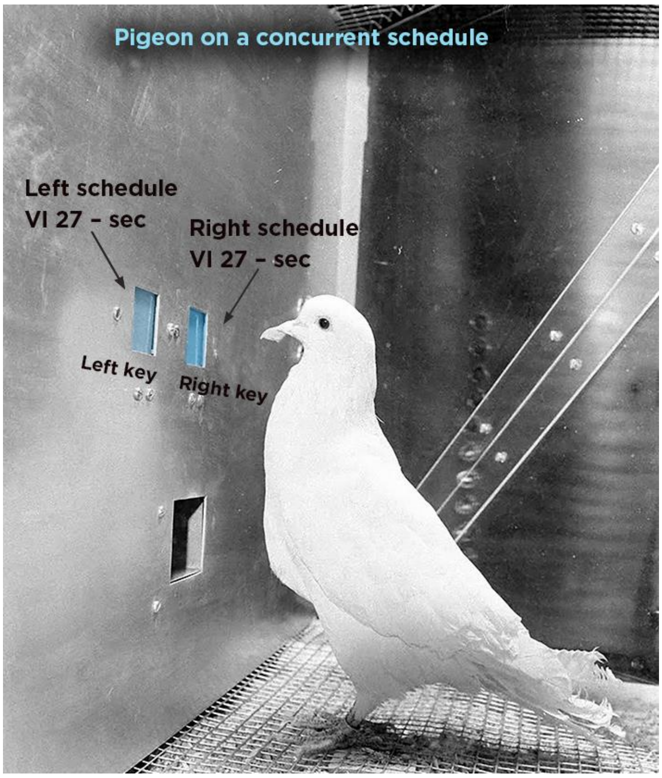

How were the specific choice behaviour datasets used to train the SNN?

The dataset used for the study presented in Chapter 4 was collected and reported by Landon and Davison (2001) (Figure 8). In their research, the authors used a *variable interval schedule of reinforcement* (VI) that was presented using a *two-key concurrent schedule procedure.*

Landon and Davison (2001) used a dual process to train six experienced pigeons numbered 61 to 66 in similar experimental procedures. First, the computer determined what VI-VI was in effect, and each session started with both keys lit in yellow. The environment for training was arranged as seven different pairs of concurrent VI schedules for each session, separated by a 10-min blackout. Each pair (VI-VI) was presented randomly without replacement from a list of seven different reinforcer ratio schemes (Table 3). Secondly, they determined when the reinforcer was delivered, which was scheduled according to a single exponential VI 27-s schedule ($p = 0.037/s$). Once a reinforcer was arranged, probabilistically, it was then allocated to one of the two alternative keys. A 2-sec changeover delay was used. They trained all pigeons under 16 consecutive experimental conditions where the range of reinforcer ratios was varied across conditions (the overall VI 27-s schedule was consistent). All events were computer generated and the responses and the time they occurred, were recorded on MED-PC software. The choice behaviour changed in each pair (VI-VI) and were examined when the data were aggregated according to the left-key over the right-key response ratios, obtained in successive *interreinforcer* intervals (i.e., the time between the last reinforcer obtained to the next). Then the logarithms ratios were calculated for the analysis.

The authors used a variety of analyses and concluded that:

- Choice behaviour, the log response ratios were adjusted to the successive reinforcer delivery, and this was directly related to the reinforcer ratio in effect in that component, and the overall range of possible reinforcer ratios.
- Multiple linear regression analysis showed that the reinforcer ratio in the prior component, exerted some residual control over behaviour in the current component, and this dissipated as reinforcers and controlled by the current reinforcer ratio, were acquired in the current component.
- In a reinforcer-to-reinforcer analysis, they demonstrated that each reinforcer moved their choice behaviour towards the alternative (key) from which it was obtained.
- When the same alternative key provided a series of consecutive reinforcers, the effect on choice behaviour became more noticeable.

In summary, they showed that individual reinforcers had regular effects on choice behaviour (Landon & Davison, 2001). However, this effect was contextual, meaning that reinforcers obtained had a directional effect on choice behaviour. Thus, the effect reinforcers had, could be affected by the key the reinforcer obtained, and the context in which it was obtained. This includes the reinforcer ratio in effect, the range of reinforcer ratios in effect, and the place in a component and sequence of reinforcers preceding it. The research added to the existing behaviour analyst research, that behaviour is orderly and changes in choice behaviours are quantifiable. Moreover, when the data collected are sufficient, the effects of individual reinforcers and sequences of reinforcers are also evident.

The pigeons' choice behaviour data from Landon and Davison (2001) formed the basis for the study presented in this Chapter (Figure 8). Initially I experimented with the datasets to decide which extraction approach would fit the datasets (B1 Appendix B, p.285). The best approach was taken and used for all manuscripts (B2, Appendix B, p.288). For this study, choice behaviour datasets were extracted from Condition 1 and Condition 6 (see 4.2.1-4.2.3). Permission was obtained to use these datasets. In addition, permission from the Journal of Experimental Analysis of Behaviour publisher was obtained to re-print this manuscript in this thesis (Appendix A).

Figure 8

A schematic representation of the choice behaviour datasets used in this thesis.

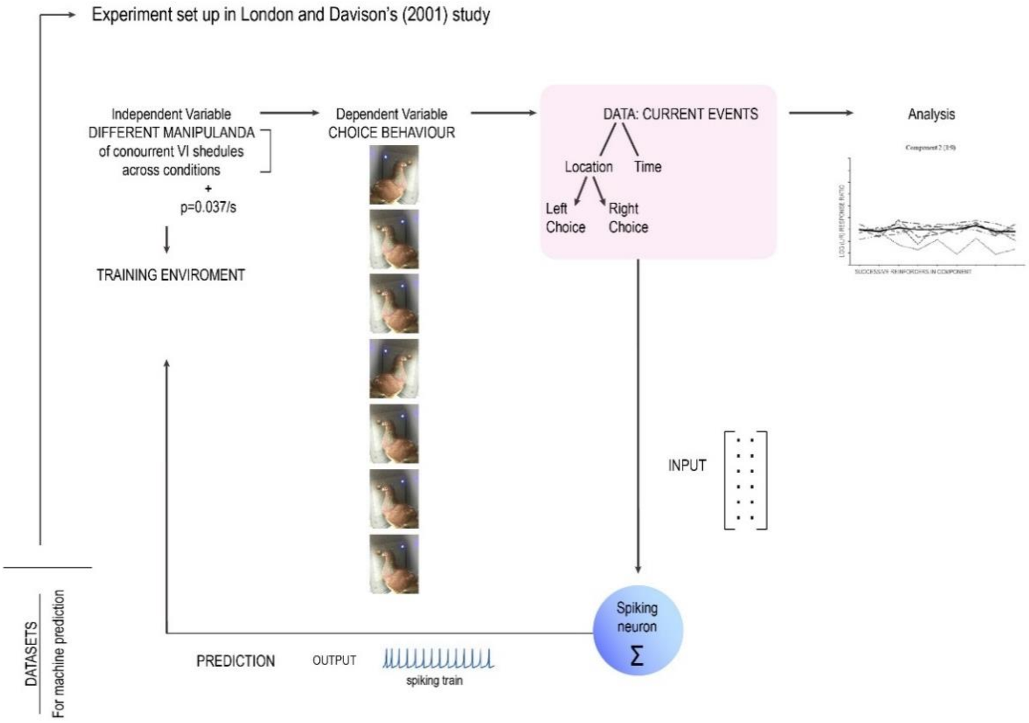

Manuscript

Plessas, Anna, Josafath I. Espinosa-Ramos, Dave Parry, Sarah Cowie, and Jason Landon. 2022. “Machine Learning with a Snapshot of Data: Spiking Neural Network ‘Predicts’ Reinforcement Histories of Pigeons’ Choice Behaviour.” Journal of the Experimental Analysis of Behaviour, April. https://doi.org/10.1002/JEAB.759.

4.1 Introduction

An analysis of the interactions between behaviour and with the environment provides us with information about reinforcement contingencies (B. F. Skinner, 1969). Reinforcement contingencies have been extensively studied in choice research and models developed to understand how behaviour changes in relation to environmental factors (Davison & McCarthy, 1988; Pierce & Epling, 1983; Skinner, 1969; Staddon, 2016). Extensive research has shown that this relationship can be complex; choice is dependent on recent reinforcement history, such that responses that have in the past produced relatively more reinforcement, occur relatively more often (Baum, 1974; Davison & Jenkins, 1985; Davison & McCarthy, 1988; Killeen, 1972). These effects are observed both when environmental contingencies remain stable over time (e.g., Baum, 1974), and also when environmental conditions change rapidly (Davison & Baum, 2003; Landon et al., 2003a; Mazur, 2016). Indeed, even when a reinforcer ratio remains in effect for just 10 reinforcers, choice varies systematically with the arranged reinforcer ratio (Davison & Baum, 2000). Thus, behaviour is highly sensitive to recent reinforcers, and functional relations learnt in the past contribute to current behaviour in a way that allows prediction of future responses under similar circumstances.

The relationship between behaviour allocation –for example, when two-keys are present– and the relative reinforcement arrangement would suggest that either can be predicted from the other. Given the fundamental relationship between choice and reinforcement, different predictive questions can be asked about the effects of variables influencing the behaviour of individual organisms. In behaviour analysis, a traditional approach would be to predict the next distribution of responses when the reinforcer ratio is set as the independent variable (Baum, 2018). This informs us on how reinforcers obtained in a particular context, over successive instances, can lead to a change in behaviour. An alternative approach to prediction is to speculate about the reinforcement history of a behaviour, that is, about the reinforcement arrangements used in training conditions and no longer in place by looking only at current behaviour. This approach allows us to identify the reinforcement history when such information is not available and inform us about the conditions in which new behaviour can be learnt. The latter bears special significance for both applied and experimental contexts. Especially in applied settings, the reinforcement history is commonly unknown, and is speculated on in the development of assessments and interventions, but otherwise, there is little attention to the matter (Pipkin & Vollmer, 2009). Different schedules of reinforcement maintain different patterns and rates of responses, and reinforcers are variables that play a role in current behaviour. However, these effects are still under investigation as earlier reinforcement history plays a role in present performance (Freeman & Lattal, 1992; Okouchi & Lattal, 2006). Recent research also supports the idea that the effects of reinforcement history are not only apparent but can reappear when the reinforcement contingency changes (Okouchi et al., 2014). Accurately identifying reinforcement history can improve our decision-making process when designing reinforcement arrangements to change behaviour. To differentiate among distinct types of predictive questions, in this paper from now on, we will use the term '*retrodiction'* to denote predicting the reinforcement history (i.e., to identify past concurrent variable-interval (VI) schedules with various reinforcer ratio arrangements).

In rapidly changing procedures, reinforcer ratios across two alternatives change multiple times in a session in an unpredictable manner. This procedure simulates natural environments to some degree, and results in behaviour that changes quickly in response to the changing contingencies (Davison & Baum, 2000, 2002; Landon & Davison, 2001). In the analysis of data from frequent changing procedures, large amounts of data are aggregated within and across sessions over a large period allowing for the analysis of extended patterns of behaviour (Baum, 2002). In this research, we wanted to investigate whether retrodiction could be possible without the need to aggregate the data and when a limited sample of behaviour was available. Therefore, we used a minimal amount of data to resemble a more limited observational data from naturalistic setting (i.e., 5-sec periods post reinforcer). In this manner, we assessed the viability of approaching observed behaviour and relevant processes in a way that might be attainable in both experimental and natural settings in an attempt assess whether a relationship is detectable using this small sample (5-s) of data.

Five second periods after the delivery of a reinforcer were chosen because of the close local proximity responses have to the reinforcer. Local effects of reinforcers, meaning the effects reinforcers have in specific locations in time, have proven to be strong and research has demonstrated individual reinforcers have short-term effects on momentary behaviour (Davison & Baum, 2002, 2003). Local effects are evident in both stable and rapidly changing environments and occur in the context of a more global shift toward the richer response alternative e.g., (Landon et al., 2002, 2003a). Localised control by reinforcers suggests that even behaviour recorded from a small interval of time should be sufficient to allow retrodiction.

Machine learning (ML) algorithms have proven to successfully complete complex tasks in only a few seconds and this advantage of ML has attracted attention across scientific domains. One subset of ML uses models inspired by neuroscience, with algorithms that simulate the properties of neurons and neural networks. Such artificial neural networks (ANN) have been used extensively to model and understand normal and abnormal brain function (Macpherson et al., 2021), as well as 'cognitive' tasks including perception and decision-making (Zador, 2019). Although ANN are widely used in healthcare (Shatte et al., 2019), education (Korkmaz & Correia, 2019b), and other fields such as speech recognition (Bhangale & Mohanaprasad, 2021) or image recognition (Cai et al., 2020), their application in behaviour analysis remains limited (Turgeon & Lanovaz, 2020). Thus, in the present study we asked to what extent ANN can identify reinforcement history based on current behaviours.

Although EAB has been hailed for its powerful experimental designs and identification of learning patterns, some argue that our scientific work has isolated itself (Poling, 2010; Vyse, 2013) from a world of modern technologies pervading almost all human activities. The aim of this investigation is to stimulate a discussion on how to approach data differently (i.e., using a snapshot of data) and rely on observed behaviour (i.e., retrodiction; direct reverse test) and relevant processes, rather than manipulating the reinforcer-behaviour relation. This novel approach can influence our decisions to determine appropriate interventions or aid knowledge over a shorter time frame. Thus, this approach might lead to dealing with behaviour in alternative ways that the behavioural community desires (Shahan, 2017).

4.1.1 Introduction to brain-inspired neural networks for this Study

ANN is a subdivision of machine learning (ML) software inspired by how the brain processes information and solves complex problems. Artificial spiking neural networks (SNNs) are termed *the third generation*[2] *ANNs,* due to their ability to provide biologically plausible neuronal models that capture some of the complex temporal dynamics of the data (see Gerstner & Kistler, 2002 for more on SNNs). A computational spiking neuron (also known as an artificial biological neuron model) produces spikes - discrete events that take place at points in time, coming from one neuron to another and connected through weights (Maass, 1997). This ability to work with discrete events that occur at separate times, corresponds well with measures of choice data. Information transfer in artificial spiking neurons, mimics the way information is transferred in the biological neuron by considering the exact time of the spike or the sequence of the spike (Figure 9). As presented in the example, the model learns by using artificial biologically plausible algorithms (Σ) feeding the information forward across the neural network, starting from the input node to the output node. The weights represent the synaptic connections of neurons and refer to the strength of those connections, i.e., the effect of the firing rates. Based on weights, the input signal may be amplified or inhibited. When the membrane potential reaches a certain threshold, the neuron will spike, generating a signal that will travel and instantly reset to a lower value. The neuron stays in a resting state for some time (the *absolute*

---

[2] ANNs can be classified into three generations according to their computational units (Mass, 1997): McCulloch-Pitts neurons, also called perceptrons or threshold gates, activation functions (sigmoid) and spiking neurons (“integrate-and-fire-neurons”).

refractory period), after which it can process new information coming from other presynaptic neurons (Gerstner & Kistler, 2002). Thus, this process attempts to resemble the natural nervous system and is considered as biologically realistic and plausible (Maass, 1997). Further, this process eliminates the need for an averaging time window and allows the processing of information in continuous time. In essence, artificial SNNs' strength is that the networks learn from spatiotemporal data.

Figure 9

A visual representation of a brain-inspired computational model

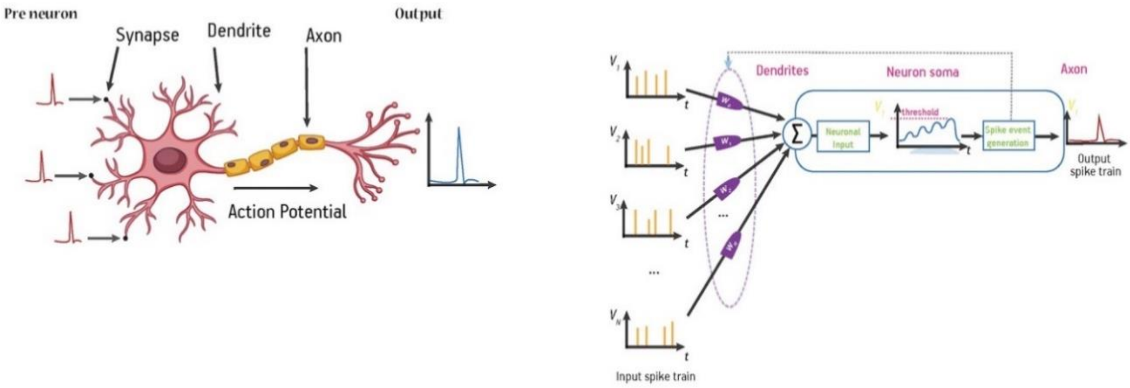

Note. On the left the graphic shows the structure of a biological neuron. On the right the graphic presents a sample of a brain-inspired computational neuron (adapted from Chen et al., 2018)

In ANN, ‘learning’ refers to the process of extracting structure from the data that will be encoded in the parameters of an artificial network and will provide all the information needed to develop the artificial network (Zador, 2019). In other words, all ANN explicitly learn from the datasets. There are three learning methods commonly used to train an artificial neural network: supervised; unsupervised; and reinforcement learning. In this study we employed a semi-supervised ML method combining supervised and unsupervised learning methods. With supervised learning the ML task is to employ the same label to the input/output set as with the training dataset. Data consists of pairs-an input item (for example, the choice made by a pigeon) and its label (e.g., the concurrent- VI schedule to the learnt behaviour). By training the ML algorithm, it searched for patterns on a 'labelled' dataset. After this training, any new inputs are cross matched with the training dataset to determine the desired output. By this, the learning algorithm can deduce a pattern by identifying a relationship between the target variable and the rest of the dataset based on the information it already has. On the other hand, with unsupervised learning, the training set contains no 'labels' in the data and the algorithm learns without needing to supervise the model (Rafi, 2021).

SNNs can be an appropriate tool for modelling behavioural data for several reasons. Firstly, choice and reinforcers are complex variables that constantly interact with each other in space and in time (Cowie & Davison, 2016a). We hypothesised that an ANN like SNN can be beneficial, since its advantage resides in handling spatiotemporal information. Secondly, SNNs lend themselves to the small-N approach that characterises work in the field; by extracting sufficient amount of data from fewer individuals (i.e., six pigeons) rather than the reverse (i.e., small amount of data accumulated across a large number of individuals) can add value to the prediction accuracy of the model when it is applied to a single organism.

Our aim was to investigate whether a machine learning algorithm capable of processing spatiotemporal information can identify reinforcement history when no previous knowledge about this history is provided. Retrodicting histories of operant behaviour may open exciting avenues in the field of behaviour analysis. From a practical standpoint, a retrodictive outcome would be possible, based solely on limited data (e.g., 5 sec periods). This could parallel situations where laboratory resources are limited or when clinicians need to decide on the appropriate intervention having only limited observational data available. Also, behaviour analysis could have an additional alternative tool in their ‘toolbox’ to study the relation between past reinforcement arrangements and current behaviour when generating knowledge.

4.2 Method

4.2.1 Dataset for training and testing the artificial Spiking Neural Networks

The dataset was extracted from Landon and Davison’s (2001) study in which a frequently changing concurrent-schedule procedure arranged seven different reinforcer ratios, in *components* (1:27, 1:9, 1:3, 1:1, 3:1, 9:1, or 27:1 – see Table 3) which changed randomly within each session. The overall reinforcer rate was constant, and each component was in effect for 10 successive reinforcer deliveries, and components were separated by 10-sec blackouts. The time and nature of all experimental events was recorded (for full details, see (Landon & Davison, 2001)). We extracted the raw data for all six pigeons in all 50 training sessions from one experimental condition (Condition 1). We then created an extraction rule to use a minimal amount of data. All choice responses made in the first 5 s after every reinforcer were extracted as a frequency event and represented as discrete events in time (i.e., temporal data). All pigeons’ choices were taken after the delivery of the first nine reinforcers within each component (a component ends with the delivery of the tenth reinforcer).

Table 3

The Relative Reinforcer Probability (Shown as Probability of Reinforcement to the Left Alternative) for Each of the Seven Concurrent VI Schedules (Referred to as Components) for Both Conditions

Component	Reinforcer Ratio (Left: Right)
1	1:27
2	1:9
3	1:3
4	1:1
5	3:1
6	9:1
7	27:1

Note. The overall probability of reinforcement per second was constant at 0.037.

4.2.2 Extraction rule of the data for SNN training

All files originally had a ".txt" extension and were generated by MED-PC® software. All files used in this study, were programmed by the original experimenters, to record all experimental events and the time at which every event occurred within the experimental conditions. To extract a snapshot of data from the whole dataset of Condition 1, the second author developed a computer code in Java. When the code was run by the first author, *samples* in ".csv" format were created, of pigeons' left- and right-preference responses following each reinforcer delivery, during a five second period. Each sample consisted of a single component and was organised of aggregate data from 10 daily sessions. This resulted in a total of five samples for each bird out of all 50 sessions. Each sample contained 90 periods of five seconds as nine reinforcers were given each day (9 * 10 daily sessions). Thus, from each pigeon we took a total of 35 (7 conditions* 5 ten daily sessions) samples, and all pigeons together generated 210 (6*35) samples. Each row in each sample represented a frequency event, which was created by calculating the ratio of left and right preferences based on the time window (5-sec) and on the actual location (L or R) where the reinforcer was delivered (see Appendix B for an example).

This approach to samples gave the opportunity to assess the explicit use of pigeons' temporal data and the actual location they occurred in time (L or R), was sufficient to learn to detect patterns without any need to add further information.

4.2.3 Dataset for further generalisation tests of the Spiking Neural Networks

For generalisation testing, we extracted three new datasets from the same pigeons, this time from a different experimental condition (Condition 6) of the original study (Landon & Davison, 2001). The six pigeons were exposed to Condition 6 for 35 sessions after the completion of experimental Conditions 1-5. The same environmental arrangements as in Condition 1 were set for Condition 6. We aggregated data from the 35 sessions in three separate ways – we first created samples each containing datasets from ten sessions (resulting in three samples per bird for each component); then we created samples of each, containing datasets from seven sessions; and finally, samples containing datasets from five sessions each. We did so in order to assess the model's dependence for generalisation on individual dataset size. Temporal data were taken, as in the training phase of the model, in the first five seconds after the delivery of the first nine reinforcers. Thus, from the ten-session samples we created dataset G1, where the data were aggregated from 90 five-second periods (9 reinforcers * 10 sessions) in each sample; from the seven-session samples we created dataset G2, where the samples contained data from 63 periods (9 reinforcers * 7 sessions); and from the five-session samples we created dataset G3, where each sample included 45 periods (9 reinforcers * 5 sessions). All in all, the G1 dataset included 126 samples (6 pigeons * 7 components * 3 samples per pigeon). An equal number (126) of samples resulted from the data extraction for G2 because the data did not generate more than three samples per pigeon. Lastly, for the G3 dataset, the number of samples was 294 due to the samples containing only 45 five-second periods each.

Table 4

Samples Extracted by the Pigeons' Temporal Data

Experimental Condition	Datasets	No of Pigeons	Points of 5-sec periods
Cond. 1	Training the model	5 Pigeons	90 periods
	Testing model	1 Pigeon	90 periods
Cond. 6	Generalisation-1	6 Pigeons	90 periods
	Generalisation-2	6 Pigeons	63 periods
	Generalisation-3	6 Pigeons	45 periods

Note: Some samples consisted of <90 points as some of the ten daily sessions ended at the prearranged time (45-minutes), and the pigeon had not consumed all reinforcers.

4.2.4 Single Spiking Neural Network architecture constructed for this study

We created an artificial SNN based on the architecture initially proposed in Vazquez & Cachón, (2010) to identify the reinforcement history (i.e., frequently changing concurrent-schedule procedure arranged seven different reinforcer ratios, in components) based on pigeon choices and the firing rates these data produced. This was achievable using the Leaky Integrate-and-Fire (LIF) model, a mathematical representation of a neuron that was trained to perform the task. The LIF model is commonly used due to its simplicity and computational efficiency, while attempting to mimic biology (see, e.g., Khadeer, 2020, for more on neuron models)

Figure 10 depicts a simplified schematic representation of the single SNN architecture (more technical details are provided in supplementary materials). The inputs were samples that represented discrete events of left and right choices based on the time window (5-sec) and on the actual location (L or R) where the reinforcer was delivered (see extraction rule of data for SNN for Training). The output of the SNN was the firing rates associated with the seven components presented in Table 3 (multi-classification problem).

These samples were transformed into a vector of values that simulates the electrical current that is injected to the LIF neuron model. Each sample produced a spike train, a sequence of ones and zeros, that was transformed into a firing rate. To achieve transformation of the spiking data, the information was encoded in the number of spikes over a specified temporal window (*rate encoding*) (see, e.g., Auge et al., 2021, for more on encoding processes). This produced a sequence of artificial spikes with a specific firing rate. Then each component was associated with a specific firing rate. It was hypothesised that the input signal produced by samples of the same component produce similar firing rates, whereas input currents of different components produce firing rates different enough to discriminate among the various components. We used JNeuCube to develop the proposed artificial single neuron model and to perform the experiments of this study. JNeuCube is a Java-based framework for building SNNs able to solve classification and prediction problems (https://github.com/Auckland-University-of-Technology/NeuCube-java). Specific parameters were set to train the artificial neuron to correctly perform the requested task (see supplementary materials, Appendix B).

Figure 10

The SNN architecture constructed for classification of seven components based on pigeons' choices

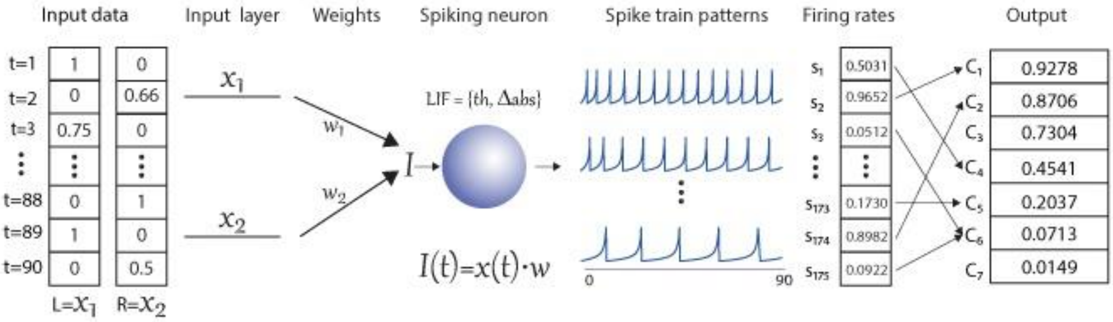

Note. The letters in the figures represent: 90 periods of five seconds as t ; left choice responses as X_1; right choice responses as X_2; left and right weights as w1-w2; s1….s175, as samples 1-175; C1..C7 as Component 1-7.

4.2.5 Training procedure of the SNN model

We chose a Differential Evolution (DE) algorithm due to its effectiveness to investigate how we can understand patterns in our data (Storn & Price, 1997). It is a kind of decision-making tool, where decisions are made to optimise one or more objectives under specific circumstances. We used a cross-validation procedure to evaluate the model’s performance by using pigeon datasets for training the model and then for validating its performance on a testing dataset that the model had not encountered before. We trained the model applying the stratified k-fold cross-validation (k=5) on the training dataset (five pigeons) and evaluate its performance using the testing dataset (the 6th pigeon). In the 5-fold cross-validation, the data were partitioned into five sets of equal sizes that were randomly used then the model was being trained using k-1 folds and estimating its retrodictive ability using 1-fold (Figure 11). The process was repeated five times, each time selecting a different fold for training (k-1) and testing, thereby creating different accuracies each time. This strategy allows for an objective, less biased and less optimistic estimation of the model's performance than other methods allow (James et al., 2013).

Figure 11

A schematic representation of k-fold cross-validation training when using right and left responses (pigeons' choices)

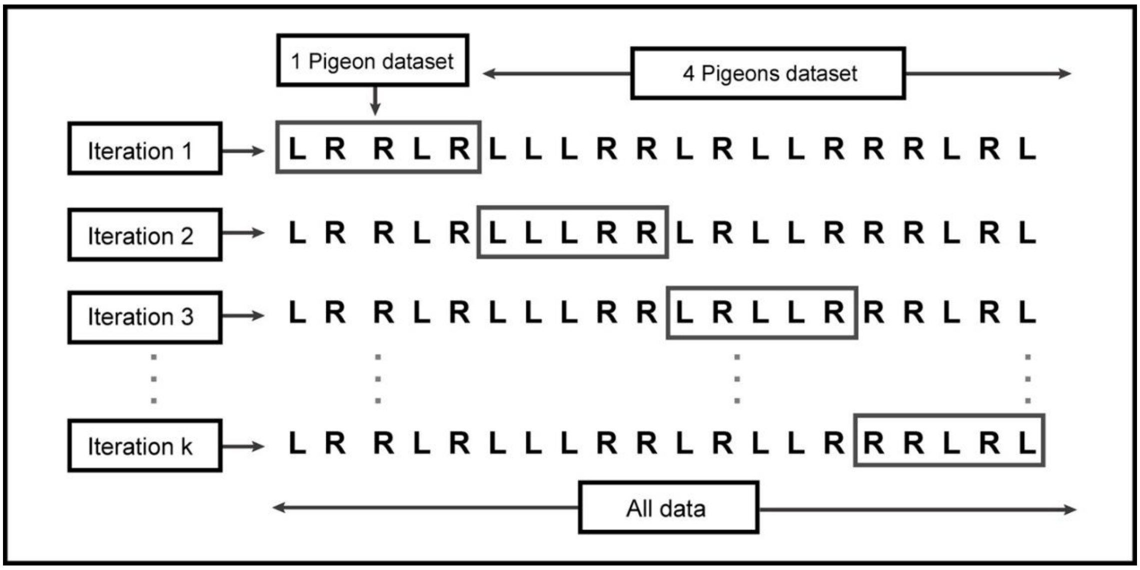

With the completion of cross-validation, we evaluated the model's generalisation ability by testing it on the data of another pigeon unknown to the model. The outcome was 40 single artificial SNNs that could accurately classify components based on a pigeon's choices. From all these, we kept the fittest as the final result. We repeated the experiment 10 times and created 60 artificial SNN models by splitting the total of six pigeons' data from Condition 1 to five pigeons' data for training and one for testing each SNN model. With this method, all possible combinations of pigeons' data were used both for training and for testing for generalisation.

4.2.6 Procedures

Model performance evaluation

We calculated five different measures to assess our model’s performance. The most common method to evaluate model performance is prediction accuracy (Table 5). However, it has been proposed that when considering using the output for clinical decisions, additional metrics should be taken into account as accuracy does not take into consideration other characteristics of the data (Kuhn & Johnson, 2018). Thus, going beyond accuracy metrics, we tested the model’s ability to identify true positives and negatives assessed by recall (or ‘sensitivity’) and specificity values. An effective predictive model that can classify components based on choice behaviour should be able to discriminate events from non-events. Responses in each component should create unique patterns that differentiate choice behaviour in one component from choice behaviour in another component. In the original study (Landon & Davison, 2001), pigeons’ responses in a component were significantly affected by the reinforcer ratio and the sequential effects of reinforcers. Having a model that performs equally well in recall and specificity is critical in terms of matching response ratios to reinforcer ratios in a component in conformity with the generalised matching law (GML). Additional metrics were also calculated to measure performance. See Table 5 for some critical notions used when measuring performance of artificial neurons.

Table 5

Performance Metrics used in this Study to Interpret Results and Formulas used for Calculations

Measure	Description	Formula
Accuracy	The fraction of correctly predicted events in relation to all data	$\frac{\#\ of\ samples\ predicted\ as\ having\ the\ event}{\#\ of\ total\ samples}$
Recall (or sensitivity)	The proportion of correctly predicted actual events (i.e.., a true positive) in reference to the total true events	$\frac{\#\ of\ samples\ predicted\ as\ having\ the\ event}{\#\ of\ samples\ with\ the\ event\ of\ interest}$
Specificity	The proportion of correctly predicted nonevents (i.e., a true negative) in reference to the total nonevents	$\frac{\#\ of\ samples\ predicted\ as\ non\ events}{\#\ of\ samples\ without\ the\ event\ of\ interest}$
Informedness	The probability of an informed decision (or Youden’s index)	$J = Recall + Specificity - 1$
Precision	The fraction of correctly predicted actual events in reference to retrieved events	$\frac{\#\ of\ samples\ predicted\ as\ \ events}{\#\ of\ samples\ predicted\ as\ true + false\ positives}$
F1	Weighted average of precision and sensitivity	$2\text{x}\frac{precision\ \text{x}\ \ recall}{precision + recall}$

Response-by-response Analysis based on the Time Window (5-sec)

We re-analysed the data from Landon and Davison (2001) study within a 5-sec window to investigate the ability of the learning algorithms to identify patterns in smaller samples by comparing the model's output with the pigeons' actual performances.

4.3 Results

4.3.1 SNN performance metrics in modelling learning histories

Analysis of the overall performance of the classification model

Overall, all combinations of pigeon datasets were able to detect the reinforcement history from current operant choices with a good degree of success. The artificial SNN models correctly identified the components based on the pigeons' choice responses, ranging from 93% (Pigeon 66 – data used for testing) to 96% (Pigeons 61 to 64) correct identification. Overall specificity performance measures were higher (≥96%) than recall, illustrating that all artificial SNN models were better at classifying which component a choice response does not appertain to (specificity) than identifying in which component choice learning took place (recall). We combined these measures to estimate informedness, i.e., error magnitude in recall and specificity, as presented in Table 6. This index had values ranging 77%-85%, indicating only small errors in the SNN model performance. The index was highest when Pigeon 61 was used for testing and dropped slightly with Pigeon 66. The precision results revealed a similar pattern of recall and informedness metrics (Table 6). These results reveal that only a few events that should have been predicted as events, were not. Lastly, F1 calculated as a weighted average of precision and recall score verified the same results. Thus, detection of patterns in choice responses with a small window is possible with SNNs. Overall, the results demonstrated high accuracy in the models' performance when making decisions regarding components and choice responses.

Table 6

The Overall Results of the Six Best Models for all Combinations

Outcomes	Pigeon 61 Overall	Pigeon 62 overall	Pigeon 63 overall	Pigeon 64 overall	Pigeon 65 overall	Pigeon 66 overall
Accuracy	0.96	0.96	0.96	0.96	0.94	0.93
Recall	0.87	0.87	0.87	0.86	0.81	0.81
Specificity	0.98	0.98	0.97	0.97	0.96	0.96
Informedness	0.85	0.85	0.84	0.83	0.77	0.77
Precision	0.87	0.87	0.87	0.86	0.81	0.81
F1	0.87	0.87	0.87	0.86	0.81	0.81

Note. The results reflect cross-validation (CV) training and generalisation testing for validation.

The results are listed per pigeon used for testing generalisation.

Analysis of the classification models' performance per component

An additional analysis was conducted to examine the models' performance in correctly detecting unique patterns in pigeon responses in each individual component. The number of errors the algorithm made in identifying the actual component for each pigeon dataset was calculated and is shown in Figure 12. The two indicative models presented in the figure show that most errors in identifying the component occurred when choice responses were trained under Component 2 (1:9 reinforcer ratio). Low recall in Component 2 was observed across pigeon datasets contrasting the classification rate for other components. Furthermore, the models' ability to identify patterns in the data was overall higher with Component 3, Component 5, and Component 7 than with Component 1, Component 2, and Component 6. Overall, these results suggest that differences in metrics presented in Table 6 were dependent on the component, as any errors in the retrodictive ability of the model were more prominent with specific components (i.e., Component 2). The phenomenon might have a specific explanation pertaining to the training environment in Component 2. The retrodictions made by our model were analysed further to investigate whether reduced performance was due to the learning algorithm, or it reflected the actual pigeons' performance.

Figure 12

The figure shows a normalised confusion matrix across all seven components with the horizontal line representing the retrodicted component and the vertical line the actual component. The diagonal elements represent the recall values. A normalised matrix represents all components as having 1.0 samples (or 100%). The seven components are abbreviated as C1 to C7

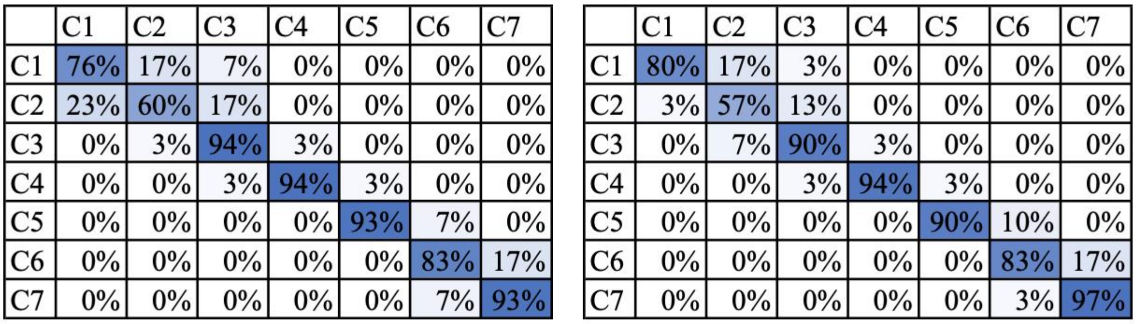

	C1	C2	C3	C4	C5	C6	C7
C1	76%	17%	7%	0%	0%	0%	0%
C2	23%	60%	17%	0%	0%	0%	0%
C3	0%	3%	94%	3%	0%	0%	0%
C4	0%	0%	3%	94%	3%	0%	0%
C5	0%	0%	0%	0%	93%	7%	0%
C6	0%	0%	0%	0%	0%	83%	17%
C7	0%	0%	0%	0%	0%	7%	93%

Trained with Pigeon 61-65, tested on Pigeon 66

	C1	C2	C3	C4	C5	C6	C7
C1	80%	17%	3%	0%	0%	0%	0%
C2	3%	57%	13%	0%	0%	0%	0%
C3	0%	7%	90%	3%	0%	0%	0%
C4	0%	0%	3%	94%	3%	0%	0%
C5	0%	0%	0%	0%	90%	10%	0%
C6	0%	0%	0%	0%	0%	83%	17%
C7	0%	0%	0%	0%	0%	3%	97%

Trained with Pigeon 62-66, tested on Pigeon 61

Note. The figure demonstrates a table that is used to describe the performance of a model (*confusion matrix*). In this study the performance of the model was evaluated by assessing the number of choice responses retrodicted as learnt in a component, and which ones were not. For example, in line 1, 76% of choices were retrodicted as learnt in Component 1, yet 17% were retrodicted as learnt in Component 2 though in Component 1 and 7% as in Component 1 though learnt in Component 1. From the confusion matrix we compute the rates of errors, true positives and negatives and so on.

Analysis of the model performance per individual pigeon dataset

We examined the effects of individual pigeon datasets on the artificial SNN model by analysing the machine metrics when splitting the metrics for training and testing data. The model handled individual datasets equally well (Table 7). This indicates that the generated models performed highly in classifying pigeon responses by component, showing high recall to the relevance of the data (>82%) and specificity

(>96%) in identifying true negatives (non-events). The precision of the model was also high (>82%), showing that the overall analysis was relevant to the whole data set.

Table 7

The Results of the Best Performing Model for all Combinations when Splitting the Data

	Pigeon 61		Pigeon 62		Pigeon 63		Pigeon 64		Pigeon 65		Pigeon 66	
	CV	**Test**	**CV**	**Test**	**CV**	**Test**	**CV**	**Test**	**CV**	**Test**	**CV**	**Test**
Accuracy	0.94	0.97	0.95	0.97	0.94	0.97	0.95	0.96	0.95	0.92	0.96	0.91
Recall	0.82	0.91	0.83	0.91	0.82	0.91	0.83	0.89	0.85	0.77	0.87	0.74
Specificity	0.96	0.98	0.97	0.98	0.96	0.98	0.97	0.98	0.97	0.95	0.97	0.94
Informedness	0.79	0.79	0.80	0.90	0.79	0.90	0.80	0.87	0.82	0.72	0.84	0.69
Precision	0.82	0.91	0.83	0.91	0.82	0.91	0.83	0.88	0.82	0.77	0.87	0.74
F1	0.83	0.90	0.83	0.91	0.82	0.91	0.83	0.88	0.85	0.77	0.87	0.74

Note. The results reflect cross-validation (CV) training and testing for validation

4.3.2 Pigeon performances compared to SNN model performances

To examine the differences in the retrospective ability of the artificial SNN model, in particular with Component 2, we re-analysed the extracted pigeon datasets (within a 5-sec window following the delivery of a reinforcer) in each component from Landon and Davison's (2001) study. For each dataset, pigeons' actual choices were aggregated according to the left-key over right-key response ratios by conducting a response-by-response analysis of pigeon's choices; logarithms of these ratios were calculated and plotted as a function of each sequential response of the pigeons. When responses were occurring only on one alternative, the log(L/R) was set as 3.5 to indicate the exclusive direction of the pigeon's choice. Our analysis revealed that birds' responses in Component 1 (1:27) and Component 2 (1:9) followed a similar pattern; choices seemed more extreme than in other components (Figure 13).

Figure 13

Log response ratios of choices emitted during the first five seconds following each successive response in each of the seven components of Condition 1 from Landon & Davison (2001).

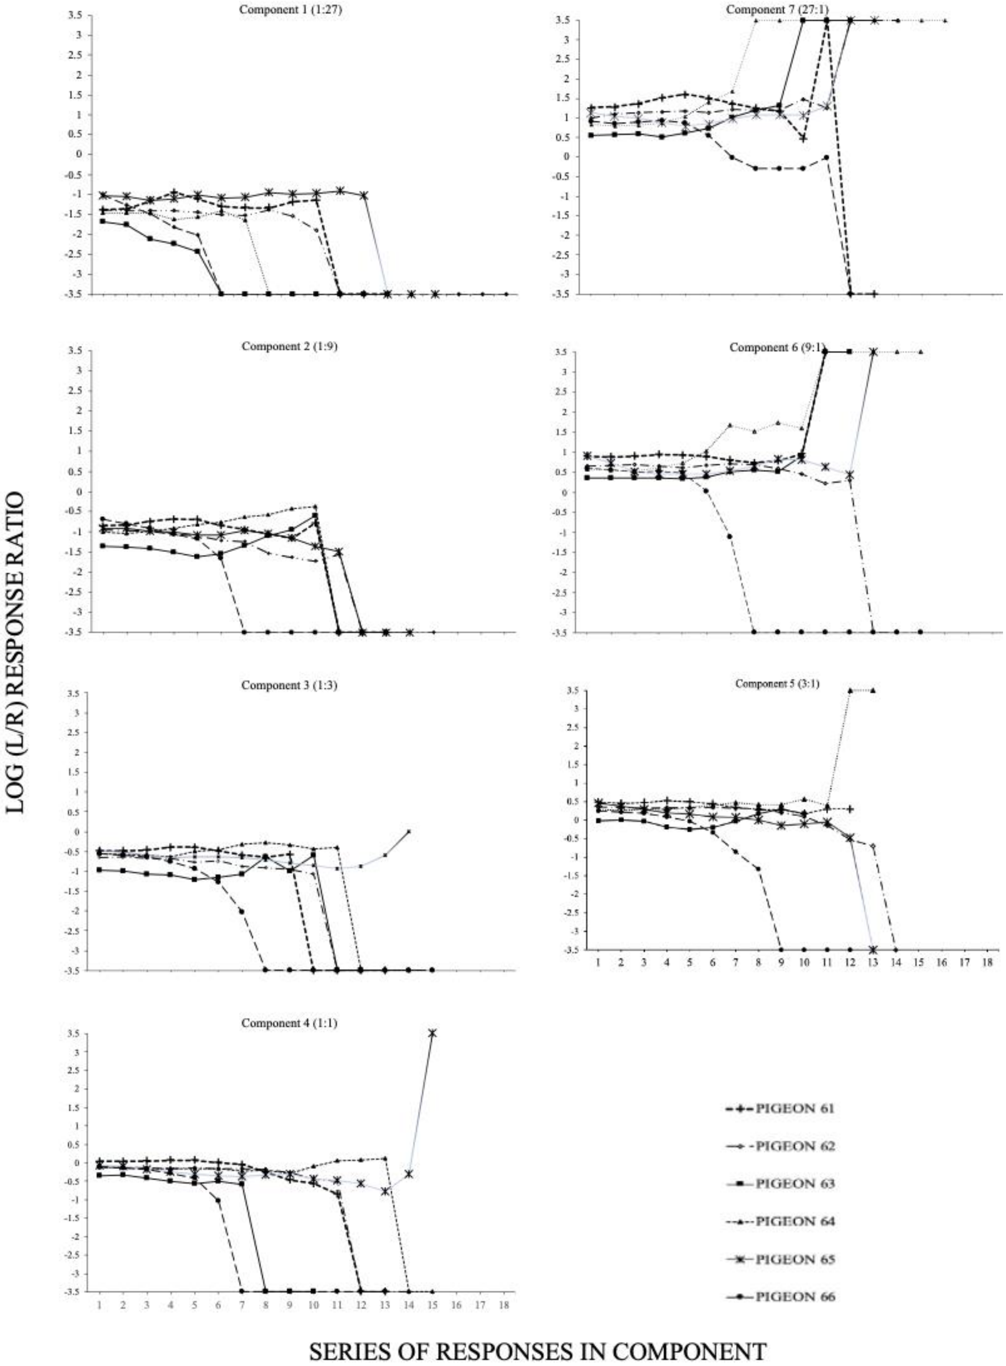

Note: Response Numbers Were Summed Across 50 Sessions. On Each Graph, Separate Plots Show Data from Each Individual Subset. On the X-axis, the Successive Responses for all Pigeons are Presented. When Responses were Occurring Only on one Alternative, the Log(L/R) was set as 3.5 to Indicate the Exclusive Direction of the Choice

Further, when we extended our analysis to all pairs of responses –as accumulated within the sequence of responding–, a similar pattern was observed (Figure 14) and became more prominent from the third response on. Across all combinations of constructed models, Component 2 was confused with Component 1 more. The analyses in Figure 13 and Figure 14 suggest that the artificial SNN model has not found a pattern because Components 1 and 2 had similar patterns; in this sense, the model reflected pigeon responses in those particular training environments.

Figure 14

Total number of pair responses emitted for the first five seconds following each successive response in each of the seven components of Condition 1 from Landon and Davison (2001).

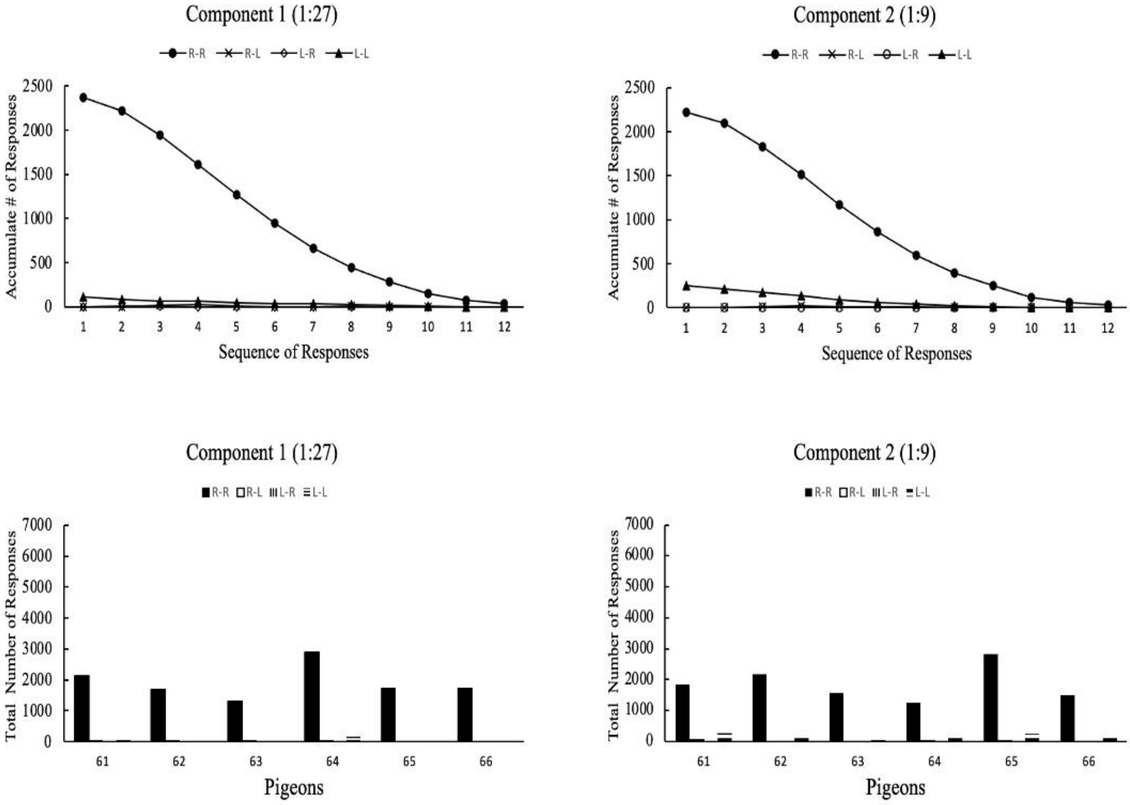

Note. Response Numbers Were Summed Across 50 Sessions as Right-Right, Right-Left, Left-Right, Left-Left Responses. On the X-Axis, the Total Number of Responses for Each Bird is Presented

4.3.3 SNN model performances in extended generalisations

Table 8 summarises the metrics for generalisation. Accuracy remained high across all pigeons (>=93%) and for all three generalisation approaches. Informedness metrics dropped with specific pigeons' datasets, which suggests that the ability of the model to identify the component was not related to the size of the samples provided but rather to the individual pigeon's' differences in patterns of responding. The only apparent exception was when tests were performed with G2, where informedness scores were slightly higher than informedness scores with G1 and G3. The same pattern is preserved in all the rest of the metrics, with a slight drop when only 50% of the data was provided (G3). Overall, the analysis for generalisation suggests that the generated models can handle new datasets well and five days of training are sufficient, when making decisions regarding components and choice responding.

Table 8

The Overall Results of the Best Model for Each Generalisation Test

Generalisation Test	#		Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
G1	10d	Accuracy	0.96	0.95	0.95	0.95	0.94	0.94
		Sensitivity	0.89	0.85	0.85	0.84	0.82	0.82
		Specificity	0.98	0.97	0.97	0.96	0.96	0.96
		Informedness	0.87	0.82	0.82	0.81	0.78	0.79
		Precision	0.89	0.85	0.85	0.84	0.82	0.82
		F1	0.89	0.85	0.85	0.84	0.82	0.82
G2	7d	Accuracy	0.97	0.95	0.95	0.95	0.95	0.95
		Sensitivity	0.89	0.86	0.87	0.84	0.86	0.86
		Specificity	0.98	0.98	0.98	0.97	0.97	0.97
		Informedness	0.87	0.83	0.83	0.81	0.83	0.83
		Precision	0.89	0.86	0.86	0.84	0.86	0.86
		F1	0.89	0.86	0.86	0.84	0.86	0.86
G3	5d	Accuracy	0.93	0.93	0.93	0.93	0.93	0.93
		Sensitivity	0.82	0.79	0.79	0.79	0.80	0.79
		Specificity	0.96	0.96	0.96	0.96	0.96	0.96
		Informedness	0.76	0.75	0.76	0.75	0.75	0.75
		Precision	0.81	0.79	0.80	0.79	0.79	0.79
		F1	0.81	0.79	0.80	0.79	0.79	0.79

Note. The results reflect generalisation tests per pigeon.

4.4 Discussion

The purpose of this study was to investigate whether an artificial SNN model could be trained to identify the reinforcement history that led to current choice behaviour based on small samples of choices in pigeon datasets. The results demonstrated that the single neuron architecture could identify learning histories by detecting learning patterns in choice responses. In particular, the SNN model could distinguish pigeon performance in one procedure (Condition 1) and then in another experienced much later on by the same pigeons' (Condition 6). Overall, the results showed that choices can be used retrospectively, to identify reinforcement history when this is unknown, demonstrating that artificial SNN models may be a useful tool for behaviour analysis.

Artificial neural networks learn solely based on datasets; therefore, they reflect patterns encoded in the parameters of the artificial network. In this study the SNN model was developed to detect patterns of two-alternative choice responses (input data), without any knowledge of what experimental conditions or reinforcement arrangements the pigeons had been exposed to. The ability of the artificial SNN model to predict (or retrodict), shows that limited observational data from naturalistic settings alone could potentially provide information on the relationship between choice and reinforcer. Further, this ability to retrodict the history of novel operant behaviours (i.e., novel as in not participating in the development of the artificial SNN model), allows us to shape future behaviours without necessarily having to know how these functional relations were established in the environment; rather, relying only on the information that the functional relation is apparent. Overall, knowing what relationship exists, might provide a particularly effective pathway when the sole interest is to set up optimal learning environments that are task-specific and learner-specific, and a small amount of current behaviour as the only source of information. Subsequently, the prediction outcome can help us decide which reinforcement arrangements can effectively shape future choices and which may not. Further, being able to determine the reinforcer history in this way allows us to investigate other matters such as how or why the reinforcer-behaviour relationship exists. The outcome obtained from the SNN model re-confirms the reality of a reinforcement-behaviour relationship regardless of the theoretical explanation for this relationship (see, e.g., Simon et al., 2020, for current debate on the mechanism of the relationship).

Our results demonstrated high overall accuracy in retrodicting reinforcement arrangements (≥93%) across all artificial SNN models. A possible advantage of using a machine learning tool compared to moment-to-moment visual analysis is that it can provide instant additional information on true and false positive and negative values. For example, confusion matrices (Figure 12) can visualise Type 1 and Type 2 errors providing us additional information on responses that were retrodicted as belonging to a component when they did not. These errors are not easily detected with visual analysis of single case data (Lanovaz et al., 2020). Therefore, the analysis approach explored here can provide more details on learning patterns which will enrich the behaviour analyst's decision-making regarding intervention characteristics.

Retrodiction was achieved quickly from a small set of robust behavioural data (spanning five seconds following the reinforcer delivery). Also, small amounts of data render artificial models simple and friendly to use, without need for expertise in training conditions. As rapidly changing environments become difficult to predict when the number of conditions increases, artificial SNN can be a useful tool to make predictions (retrodictions) in an efficient manner. This is of interest to experimental and applied researchers because collecting continuous data for extended periods can be tedious and costly and by using ML, we demonstrate that the first 5 second periods may not be a constraint. A further question is whether less or more than 5 second of data can alter the model's ability to detect the reinforcement history.

The literature has shown that learning occurs over time and behaviour takes time to stabilise. Thus, an additional factor to consider is how large a 'window' of data we choose to model the reinforcement history. It may also be important to consider the time needed for a pigeon to learn in a reinforcement arrangement, as both short-term and long-term relations between choice and reinforcement are evident (Landon

& Davison, 2001). Even so, moment-to-moment analysis with 5-second periods (Figure 13 and Figure 14) revealed that the pigeons behaved similarly in Components 1(1:27) and 2 (1:9) despite different learning histories in the two reinforcement arrangements. However, this was not the case for Component 6 (9:1) and 7 (27:1). A second question is from which point in time (e.g., the beginning, the middle of the training or when behaviour stabilises) we extract our data from to create inputs for the SNN to make an accurate retrodiction.

Based on research into local effects of reinforcers, we included all responses immediately after the reinforcer delivery, both from training and stabilised learning. The point in time at which training occurs in the animal's learning history may affect the ML model's ability to retrodict histories correctly. If behaviour changes with experience, then models developed with data both from early (learning) and later (stable) sessions are trained to detect patterns from variable data, which adds to the complexity of 'retrodicting'. By including all types of data as inputs we showed that the artificial SNN models could handle variability over time and modelling histories based on pigeon responses without aggregating or averaging the responses (something that is commonly required with second-generation ANN (Alaloul & Qureshi, 2020). Choice research has shown that previous reinforcer deliveries affect subsequent behaviour, and the dynamics of this relationship is determined by environmental variation (Davison & Baum, 2000). We considered that, if a series of cross-sectional data were taken (i.e., data from a set period of time during the experiment), we would have missed significant information about the degree of environmental variability. In this study the pigeons' datasets of left and right-key responses were altered by the weights of the model and combined into one element to produce one current (feature), which we used to stimulate the artificial neuron model and produce sequences of spikes. This total transformation of the data allowed us to make reliable retrodictions without the need for curve-fitting the data (Davison & Elliffe, 2009).

Yet, reducing the sample sizes (five sessions of training versus ten) reduced only slightly the model's ability to retrodict, implying that at least the sample size may play a small role below a certain number of time-events. Future studies could explore the relation between prediction accuracy, where in time the data are extracted from, how much time from the delivery of the reinforcer, and the sample size we generate as input data. Understanding how these factors impact the model's ability to predict may provide further insight into the relationship reinforcer- behaviour, and the effects of reinforcers on behavioural changes with experience. Overall, we demonstrated that artificial SNN allow us to investigate how additional environmental variables such as time in training can detect reinforcement histories with minimal amount of operant choice.

The focus of our study was on the output of an ML model demonstrating retrodiction of reinforcement histories from current behaviour when these histories are unknown, and how this outcome can be used to achieve insights into learning. Future work can take further steps to generate new research questions and investigate how a dataset is classified by alternative ANN models to compare the information derived from other models' performance. Such a comparative study of different ANNs could investigate how inputs are handled differently by other architectures and algorithms. It is also interesting to compare ML results with the human expert classification of the same datasets from experiments arranging high- and low-discriminability histories. Such a comparison might highlight what the specific abilities are that ML can contribute to behavioural data processing.

By being able to identify the most likely learning histories, this tool has the potential to be used in two separate ways: Firstly, in the experimental analysis of behaviour research to further investigate fundamental behavioural principles, and secondly, to understand what maintains current behaviour and to identify effective training conditions that can produce the desired future behaviour. The approach we used proves to have a good utility in experimental contexts with simple choice paradigms (such as left- vs. right-key responses). Future studies can extend to more complex choice situations, as ML can be accurate with complex data. If more inputs are provided (e.g., three-key alternative responses, or measures of inherent bias), it is not clear yet whether the prediction ability of the model would remain high or even improve in situations where responses follow a similar pattern but were trained under different histories (as with Components 1 and 2). So far, with the limited information provided to our model (left- and right-key responses within a 5-sec window), the SNN approach was effective in identifying reinforcement histories. Therefore, an artificial SNN can be used with datasets consisting of behavioural responses alone; in other words, it does not require the inclusion of other data from the learning history and its use is simple and accessible without having to implement common training-testing procedures.

The performance metrics of the artificial SNN model can inform the development of new hypotheses by revisiting the actual datasets. Artificial SNN can complement the existing understanding of reinforcer-behaviour interactions. These algorithms allow us to pose new research questions and detect relationships that shed light on the mechanisms of choice behaviour of individual organisms. For example, the analysis of the moment-to-moment behaviour revealed that choices in Components 1 and 2 showed a similar pattern, with response initiation being quicker when the rich key was on the right. Moreover, we saw that, when the rich key was on the left (Components 5 to 7), responses often occurred on the right key where no reinforcement was arranged. The pigeons tended to respond to the right key more than expected, regardless of the arranged reinforcement, reflecting a right-key bias (Baum, 1974). This bias meant that choice in left- versus right-key conditions was not symmetrical, which may have hindered the model's training and its resulting ability to differentiate learning patterns, if any. This factor may be critical to making a pattern discernible with some histories, i.e., when biases are present, or reinforcer ratios are similar. Nevertheless, it is worth reiterating that even with this bias present, the artificial SNN model was able to detect learning histories well using a small amount of data. This indicates that SNN offers an advantage beyond conventional approaches. Future research could investigate how biases and other sources of 'confusion' (e.g., smaller ranges of reinforcer ratios, contradictory signals) may impact on the ability of SNNs to accurately identify learning histories.

Artificial spiking neural network modelling has emerged out of an interest in modelling the behaviour of the biological neuron to understand human behaviour. Its ability to provide insights into patterns of the brain by analysing the effect of environmental stimuli on the spatiotemporal brain data has led to substantial research interest in modelling brain with artificial SNN (Ghosh-Dastidar & Adeli, 2009); i.e., if we model how the brain works, we could understand the behaviour of a living organism (in-out approach). Here, we reversed this direction by analysing spatiotemporal information exclusively from reinforcement arrangements under which a response occurred, thus detecting patterns that are also helpful in understanding a living organism's behaviour (out-in approach). Contemporaneous computerised machine learning tools like artificial SNN can open avenues to more complete accounts of behaviour. This study, to our knowledge, is the first to use direct measures of behaviour (choice responses) with an artificial SNN model and illustrates how analysis of within-subject designs and small group participants can be used to answer alternative questions (i.e., by training and testing a model, what behavioural data can tell us about learning histories rather than vice versa).

The modelling here has advanced our knowledge beyond what is known from traditional analyses: we found that snapshots of data from current learnt behaviour and, importantly, data that are variable and extracted from unpredictable environments, contain patterns detectable by our ML model. Experimental analysis of behaviour has focused mostly on current behaviour-reinforcement contingencies. SNN modelling's special contribution can be its capability to transform the temporal data for us to analyse both from current contingencies and from past experience extremely fast and by using the same algorithm. There is a need to focus on past reinforcement history (Freeman & Lattal, 1992; Okouchi et al., 2014), and our reverse engineering approach can augment existing methods of looking into this. The results are promising as they illustrate how ML modelling can have a translational ability when using data from highly controlled conditions to answer applied questions. Even without us knowing how these functional relations were established in the environment (which is also an essential experimental question to respond to), ML detects past learning and gives us a basis to anticipate how future behaviour is shaped. Thus, data from experimental studies can acquire clinical utility. The results promise to open some helpful avenues for translational research as we can now make use of this additional means of investigating learning patterns.

4.4.1 Conclusions

This study shows how a novel machine learning (ML) tool can inform us about behaviour-environment contingencies using a small ‘window’ of data without using the common testing-training procedures. It also demonstrates how an ML tool can be utilised as a hypothesis generator directing us to look into specific behavioural data for further investigation and analysis. Artificial SNN is a new subdivision of neural networks that we preferred over other ANNs, because SNN can process spatiotemporal behavioural data, an ability that makes SNN a promising tool given the role that elapsed time plays in changes of preferences. In sum, we found that artificial SNN allowed us to identify learning histories of current learnt behaviour within seconds, from datasets of organisms previously unknown to the machine. Moreover, the results confirm previous research findings about the critical impact learning history has on current behaviour. ANN can arguably contribute to laying the foundations for new training methodologies using optimal training conditions for specific learning goals and for specific individual organisms. This knowledge (that a relationship exists rather than why it exists) also re-confirms the reinforcement-behaviour relationship regardless of the theoretical explanation of this relationship. This becomes possible with the ability of ML for prediction (or ‘retrodiction’) of learning patterns. It was simple and not time-consuming to train our SNN model and more research could further optimise the model’s training and efficiency with small datasets.

An interesting future perspective is how artificial neuronal modelling will perform in classifying behaviour of neurologically more complex organisms such as mammals. If the model’s performance is comparable to its retrodictive ability with pigeons, that would indicate its ability to reliably detect choice patterns in changing environments. Such identification of response patterns could then provide a useful indication of brain processes underpinning learning, a subject matter of current neuroscientific research.

4.5 Skills acquired as a researcher

Before and during this first experiment, I participated in three AI postgraduate papers with AUT (NeuroInformatics, Data Mining and Artificial Intelligence/ Ethics in AI). My participation in classes and reading literature helped me to acquire a good understanding of machine learning methods. As a result, all areas mentioned below were enriched and strengthened for the duration of this doctoral thesis:

- Introduction to Python coding.
- Introduction running SNN experiments and comparing the results with traditional ML models.
- In-depth understanding of choice behaviour and pattern-detection

4.6 Links between Chapter 4 and Chapter 5

In Chapter 4, a simple SNN containing just one LIF neuron could handle binary choices as inputs, to retrodict their unknown learning histories. The approach was novel, and therefore, the outcome of this experiment determined the next research steps. The novelty was in the dataset used, featuring small datasets of a 5-sec window of binary choices with the reverse-testing procedure applied. Table 2 (on page 64) provides a summary of contributions.

The next step was to validate its effectiveness by comparing the performance with other types of ANNs. Therefore, the same datasets as in Chapter 4 were used to train the SNN along with other deep learning models. In Chapter 5, we evaluate the outcome obtained from the single-neuron SNN that receives inputs from event-based data to other ANNs that sequence the data. Therefore, this comparison provided more robust validation for the single-neuron SNN performance.

Chapter 5 Performance comparison of spiking neural networks to other neural networks in predicting reinforcement history of pigeons' binary choices

Prelude

The study in this Chapter compared the single-neuron SNN in Chapter 4 with two deep-learning neural networks – a recurrent neural network (RNN) and a convolutional neural network (CNN). These two neural networks were chosen due to their flexibility during training and because they have been tested for various tasks. Though CNNs are renowned for image processing and RNNs for time series data, both have strong evidence for dealing with sequential data (DiPietro & Hager, 2020; Kiranyaz et al., 2021). Sequential datasets would require that each point in the dataset is dependent on the other point that follows. This includes for example, video data, audio data, text and so on. AI tools that can process sequential data were considered as the best alternative when successive binary choices were used as input data, as it pertains to the orderly effect of data in time.

Furthermore, RNNs and CNNs are engineered to handle data differently, which can influence how these models learn during training. In this Chapter, we wanted to see if these differences had an impact on the predictive outcome for each of these DNN models and how outcomes compared to the single-neuron SNN.

At the commencement of this Chapter, there is an opening section with information on the deep learning models used in the paper, presented in this Chapter. This information was not included in the manuscript and is included here to provide a basic overview of the architecture of DNN models.

Recurrent Neural Networks

RNNs have proven useful with predictive tasks that rely on a sequential input, which is data presented in a sequence, such as the words in a sentence, to process the information. The essential function in an RNN is mapping the input to an output sequence while taking each step of previous information into the loop, thus providing an alternative to a feedforward structure for a neural network (e.g., SNNs) to learn. Therefore, RNNs are engineered to create dependencies between previous and current information to identify more complex patterns. However, the memory in RNNs is a condensed representation of the entire sequence preceding the current data. This condensing occurs at every step, meaning that data at the start of the series has been condensed multiple times. As the sequence grows, the information at the beginning of the series is slowly degrading, making it difficult to create long-term dependencies.(Bengio et al., 1994; Boné & Cardot, 2005). Thus, if the data sequence is very long, the neuron may have difficulty predicting the final output. One way to overcome this shortcoming was to develop two alternative recurrent units, or activation units: the *long short-term memory* (LSTM) and the *gated recurrent unit* (GRU). The LSTM uses internal memory, the *cell state*, where new information is added, and old information is *forgotten*. The amount of information that remains in memory or is forgotten is decided by the *forget gate* (a mathematical function), and the amount of new memory content added to the memory cell is determined by the *input gate* (Hochreiter & Schmidhuber, 1997). Thus, the LSTM is able to decide whether to keep memory content through the gates instead of replacing the content of each time-step. This architecture allows critical features in the data to be carried on over longer periods and capture long-term dependencies in the data (Graves, 2012). GRU is a simplified version of the LSTM architecture. While LSTM has gate functions and distinct memory cells to control the flow of information, GRU takes a simpler approach by not separating the memory cells. Like LSTM, GRU also utilises gate functions, but it keeps things simpler by not separating the memory cells. Despite this simplicity, GRU remains equally effective and productive (Chung et al., 2014).

Convolutional Neural Networks

CNNs have shown significant capabilities in image processing and learning to self-optimise and better differentiate one image from the other (Li et al., 2022). All standard techniques used to improve ANNs also apply to CNNs. What distinguishes CNNs from other ANNs is the *convolutional layer*.

Figure 15

An example of the convolutional process and max pooling using a 3x3 filter (kernel).

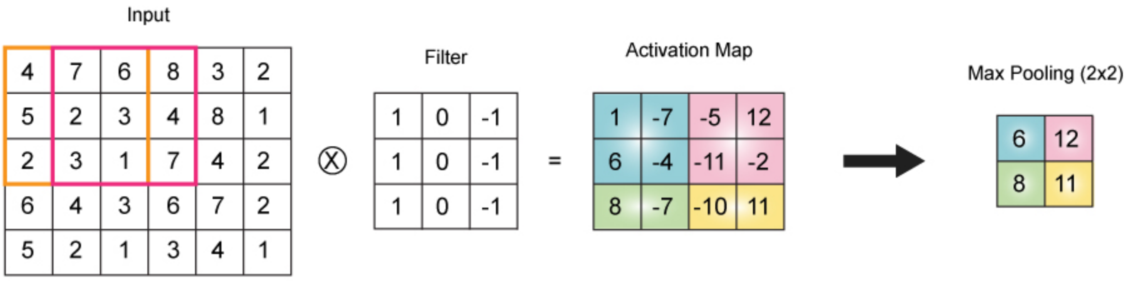

This consists of parameters denoted as *filters* (or *kernels,* small grids of cells as in Fig x) applied throughout the training phase of the CNN. Each layer *convolves* the input image and creates an *activation map,* which is a 3-D map of features, including height, width and colour. During training, each filter is designed to look for a specific pattern. When processing the data, these filters are applied across the input, and the level of activation in the corresponding activation map is determined based on the extent to which the input matches the learnt pattern. In other words, if a particular area of the data closely aligns with the pattern recognised by the filter, a higher activation will occur in that specific part of the activation map.

As represented in Figure 15, the filter slides all along the image, multiplying the value of each filter cell with the value of the underlying image pixel. Each neuron processes one subregion of the image. Subregions may overlap, therefore there are connections between the neurons. Though the neurons share parameters, every neuron has its own input and an assigned operation. Every component of the activation map acts as the output of each artificial neuron connected to a specific subregion of the input image (local connectivity). The activation map is stacked for every kernel to produce the convolutional layer output. A *max-pooling layer* further reduces the number of parameters by *down-sampling* the spatial dimensionality of outputs, and the maximum value in each subregion is the output of the filter. Thus, the last layer contains a loss function associated with the predictive outcomes (i.e., class labels) (Mostafa & Wu, 2021; O'Shea & Nash, 2015). The output of the max pooling is the new activation map (Figure 15). Therefore, training the CNN includes the optimisation of filter cell values in convolutional layers, as well as of the weights in the last phase of the fully-connected layer, the last layer in the hidden CNN.

Manuscript

Anna Plessas, Joshua Bensemann, Josafath I. Espinosa-Ramos, Sarah Cowie, Jason Landon, Dave Parry. "Performance comparison of spiking neural networks to other neural networks in predicting reinforcement history of pigeons' binary choices", 21 March 2023 submitted for review at Neural Computing and Applications.

5.1 Introduction

Recent work by the current authors has used machine learning (ML) to explore predictive relationships between known learnt behaviour and unknown training environments (Plessas et al., 2022). However, whether this predictive ability remains in effect when deep learning techniques are used, is still unanswered.

Behaviour is a naturally occurring phenomenon with fundamental properties that comply with the laws of physics, as any other natural phenomenon. In behavioural psychology, to understand how learning develops, behavioural experiments have often focused on the study of choice behaviour and, in particular, on what factors affect choice between behaviours over time. These experiments reveal how environment affects behaviour, and thus how environments may be changed to precipitate behaviour change. Every organism's response, at any given time, is discerned as a choice behaviour (e.g., a peck to left and right keys being two separate behaviours a subject can choose between), and under this conceptual framework, research has focused on understanding the interplay between competing rewards and the power of switching contingencies from favouring one response (e.g., pigeon pecking on the left-key) to an alternative one (Herrnstein, 1970). Research has shown that animal choice conforms to the same general rules as human choice (Cowie et al., 2021; Kelley et al.,

2017) and can assist our decisions in determining appropriate therapeutic interventions.

Choice is often recorded as a datum of relative rates of responding. For example, in an experimental chamber, we heavily *reinforce* by rewarding, a pigeon's choices for the left key rather than the right key, and the pigeon will peck on the left key in at a much higher rate, thus showing how relative rates indicates adjustments in learning (adaptation to environment).

Predicting which choice an organism will make requires the experimenter to:

a) design how rewards are given for each key,

b) observe how the participant shifts its preference from one key to the other after the reward delivery,

c) decide what method to use to quantify behaviour as the chosen method can influence conclusions (e.g., Boutros, Elliffe and Davison, 2010). Only then, can we make inferences about what will happen under similar conditions in the future. This methodological approach enables specific predictions to take place and the mathematical models developed to quantify choice behaviour are designed to explore how environmental variables affect behaviour (Baum, 1974; Davison & McCarthy, 1988). Therefore, these are primarily descriptive, and less able to make all sorts of predictions. For example, the existing quantitative models, are less able to predict explicit experience (i.e., past training environments) that shaped a particular behaviour that occurs at present. Being able to make predictions in behavioural research is important, and innovative technologies can open avenues to answering novel questions that enrich our understanding of learning mechanisms beyond interpretive accounts. The technological ability to answer diverse predictive questions minimises the gap between laboratory and clinical practice, as the laboratory will have the tools to deal with phenomena that have been ignored (Killeen, 2018) such as patterns in behaviour that have gone undetected so far. Consequently, by default, our current predictive ability is constrained by the mathematical models available, and new tools can assist us to develop deeper conceptual interpretations of learning processes by allowing reverse testing to be conducted. A novel approach, such as machine learning algorithms, can contribute to the advancement of behaviour research through new theoretical insights founded on objective data (Vyse, 2013).

A recent study (Plessas et al., 2022) used a single-neuron spiking neural network (SNN) to investigate whether it could retrodict pigeons’ unknown reinforcement history, (i.e., predict the past), if pigeons’ binary choices were used as input vectors. The approach was novel in that we conducted a reverse test with the use of SNN and binary choices shaped across time. The authors trained the SNN on previously published behaviour datasets and asked to what extent a machine could identify the reinforcement history by detecting patterns in pigeons' choices (associated with specific reinforcement histories). They trained the model using a Leaky Integrate-and-Fire (LIF) neuron as the temporal activation unit and used k-fold cross-validation with a differential evolution algorithm to classify and optimise the LIF threshold and refractory time. Overall performance across all developed models was high (Accuracy scores> 93%; F1 scores>81%), indicating that spiking neural networks can make predictions exclusively from behavioural observational data (such as a choice made). The SNN models’ predictions were compared to pigeons’ performances using a snapshot of data (i.e., binary choices and their reinforcement history) and the comparison revealed that pigeons’ responses did follow the pattern reflected in the models' outcome.

It was hypothesised that event-based processing used in the abovementioned study might be suitable for choice behaviour data processing because it allows outputs to be derived from spatiotemporal information (Gerstner & Kistler, 2002) and discover patterns without aggregating or averaging the data. SNNs are considered the third generation of ANNs and are energy-efficient, promising models that mimic a biological brain's neuronal dynamics (Maass, 1997; Taherkhani et al., 2020). Choice research has illustrated that control, and therefore prediction, is an outcome of time and location of the reinforcer, as well as the animal’s ability to discriminate between stimulus-response and response-reinforcer interactions (e.g., Cowie and Davison, 2020; Davison and Cowie, 2022). Thus, choice behaviour and reinforcers are complex variables that constantly interact with each other in space and time (e.g., Cowie & Davison, 2016a). The effects of reinforcers differ in several ways depending on the context in which they are obtained (i.e., reinforcement arrangements), their location (i.e., left- or right-key alternative), and timing (e.g., after successive rewards for a particular alternative choice made) (Landon et al., 2003a; Landon & Davison, 2001). The time elapsed since a given choice behaviour was trained via reinforcer arrangement also plays a crucial role in preference change (e.g., Krägeloh, Davison and Elliffe, 2005). Choice research has demonstrated that reinforcer control is strong when choice occurs close to the reinforcer delivery (Davison & Baum, 2000). However, if changes occur long after the reinforcer was delivered (or if more changes occur in the time elapsed before a reinforcer is available again), the detected reinforcer differential becomes weaker. The weaker the apparent differential, the weaker the control becomes (Cowie et al., 2014; Miranda-Dukoski et al., 2014). Therefore, the data we feed into our artificial model can interfere with the machine’s ability to predict outcomes, and we need to be cautious about the complexity of our data leading to overestimated results or methodological inconsistencies (Cearns et al., 2019).

In this study, we wanted to evaluate whether the results obtained from a single-neuron SNN (Plessas et al., 2022) were replicable by using second-generation neural networks and comparing their performances to investigate whether their predictive ability is comparable with deep learning techniques. We used two types of recurrent neural networks (RNN) and a convolutional neural network (CNN) as a tool to explore if they could be used with temporal data such as pigeons' choices and, how they would perform relative to the results obtained from the single-neuron SNN.

The difference between RNN and other neural networks is that it takes information from prior inputs to influence the current input-output to classify time series data. Therefore, the output is dependent on the sequence. Though the choice data used were discrete events occurring in time, we thought that RNN could handle these orderly data, allowing the data to be processed sequentially. We used two types of RNN, the Long short-term memory (Hochreiter & Schmidhuber, 1997) and the Gated Recurrent Unit (Chung et al., 2014). CNN, on the other hand, is a well-established deep learning model commonly used for images. It uses convolutional layers, which contain filters that facilitate learning spatial hierarchies of features (LeCun et al., 1998b). CNNs are typically used in image processing but are also popular for classification tasks using sequential data.

Output decisions of a constructed model rely on a two-fold process- the input data and the algorithm selected by the experimenter to detect relationships, on one hand, and on the machine's automatic ability to learn, on the other hand. By using different learning algorithms and methods to transform our data, we can validate our results further and provide additional tools for understanding behavioural data from behavioural experiments. This may also allow laboratory work to answer questions of interest to the clinical practice and provide a means for minimising the gap between the two. Further, the comparison of neural network models can provide insights into how behavioural data can be used with AI and help researchers identify the strengths or weakness of each neural network when predicting learning histories from current behaviour (e.g., ability to generalise, ability to improve performance and the like) and therefore, help behavioural psychologists to understand to assess if a model is effective enough or needs improvement.

5.2 Method

5.2.1 Data

In behavioural psychology, a reinforcement arrangement is a procedure that defines how the environment will respond to the behaviour. The reinforcer is a stimulus given after the exhibition of behaviour to change the occurrence of that same behaviour in the future; in experimental settings, by *reinforcer,* we mean the food provided to the participant and by *arrangement* the rule used to deliver the reinforcer (Staddon & Cerutti, 2003). The pigeons used in this study were trained to shift their choice allocation in an experimental chamber when: each choice was linked to a specific reinforcement arrangement and choices made were observed (for more on pigeons' training see Landon and Davison, 2001). The dataset was the same as in Plessas et al. (2022); samples were created by extracting pigeons’ binary choices from one experimental condition, during a five-second period following the delivery of the reinforcer. Written consent was obtained by the authors to use the datasets. Using the same datasets with the mentioned study, direct comparisons of the performance of SNN, RNNs, and CNN were feasible.

Each sample was converted into a matrix with 90 rows (i.e., 90 five-second period events) and two columns (i.e., the ratio of left and right choices in five-second periods). The columns represented Left (L) and Right (R), whereas the rows represented the sequential data, i.e., t0 (row 1) = data after reinforcer 1, t1 (row 2) data after reinforcer 2 and so on. There were 90 rows because each sample represents ten consecutive experimental sessions, and data were extracted after each session's first nine reinforcers. The labels for each sample were one-hot encoded, resulting in a vector with seven elements. One-hot encoding means that one element is set as 1 and all others as 0. The position of the 1 represents the original value of the label, e.g., [1, 0, 0, 0, 0, 0, 0] = 1, [0, 1, 0, 0, 0, 0, 0] = 2. This encoding is useful for training neural networks to classify data.

The samples were pre-processed using standard scaling,

$$z = (x - u)/s$$

where the final output *z* is the result of subtracting the mean of all samples *u* from the original sample *x* and then dividing by the standard deviation of all samples *s*. This calculation was applied column-wise, meaning that L data were only scaled according to L data and R were scaled according to L. Pre-processing is standard for neural networks, as the models perform best when all features (i.e., L and L) are on the same scale, as is true for many statistical models. Min/Max scaling was also tried but did not impact the overall performance of the trained models. Figure 16 presents the input for all neural networks for this study.

Figure 16

Schematic representation of input data used for neural network.

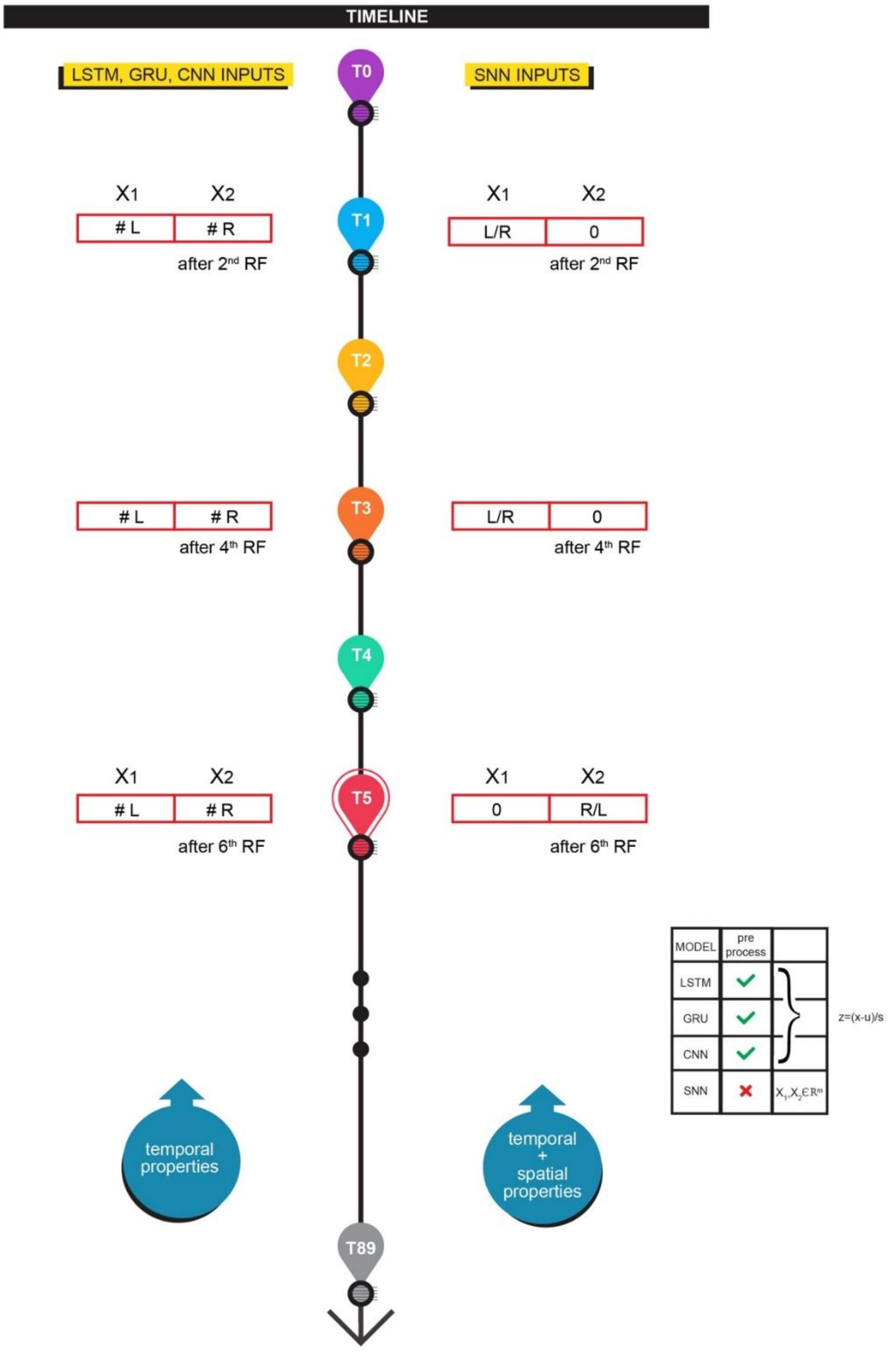

Keys: T=5-sec periods after the reinforcer delivery, X_1,X_2= responses L=Left R=Right

5.2.2 Neural networks architecture

Two types of deep neural networks (DNNs) were used to predict the reinforcement arrangement (i.e., the environment that resulted in the current behaviour) from the data: recurrent neural networks (RNNs) and a convolutional neural network (CNN). Unlike other predictive models, artificial neural networks can have as many hidden layers and learnt parameters as desired. We restricted the RNN and CNN models' size to a single hidden layer to keep the models simple (Figure 17).

Figure 17

Schematic representation of a deep neural network with a single hidden layer.

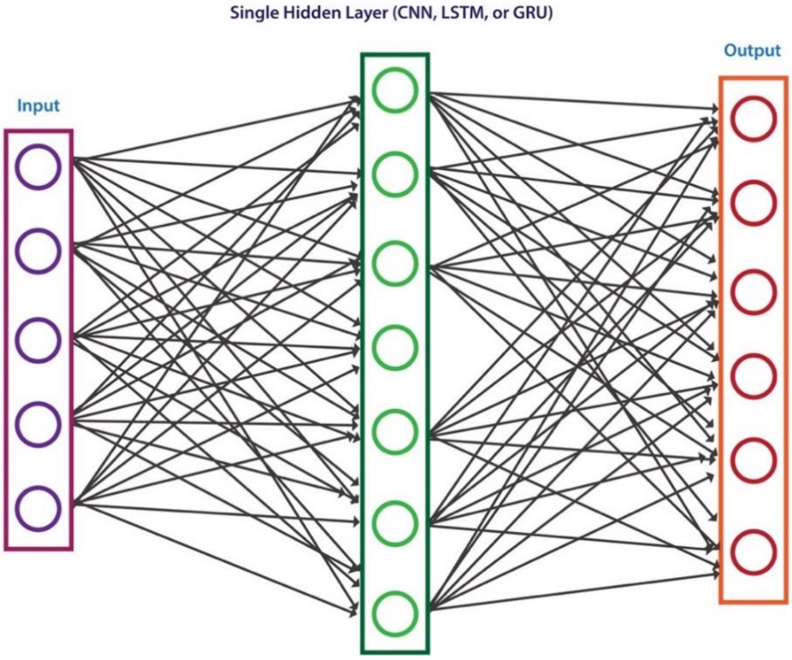

Recurrent Neural Network Models

Recurrent neural network models are typically used for sequential data, where information from time step t is retained and passed to step t+1, so that the past partially informs each step. We used two different RNN models; one used LSTM layers and the other used GRU (Figure 18). Both layer types contain gates that determine how much information is passed to the following part of the sequence. LSTM units have three gates, an input gate, an output gate, and a forget gate. The forget gate filters out old information that is no longer important. The input gate determines which information is essential for the next time step. The output gate modifies the output of the current time step. GRU units are simplified versions of LSTM units as they only have an input gate and a forget gate. The output of those units was input into the final layer.

Figure 18

Schematic representation of the recurrent neural network with a single hidden layer.

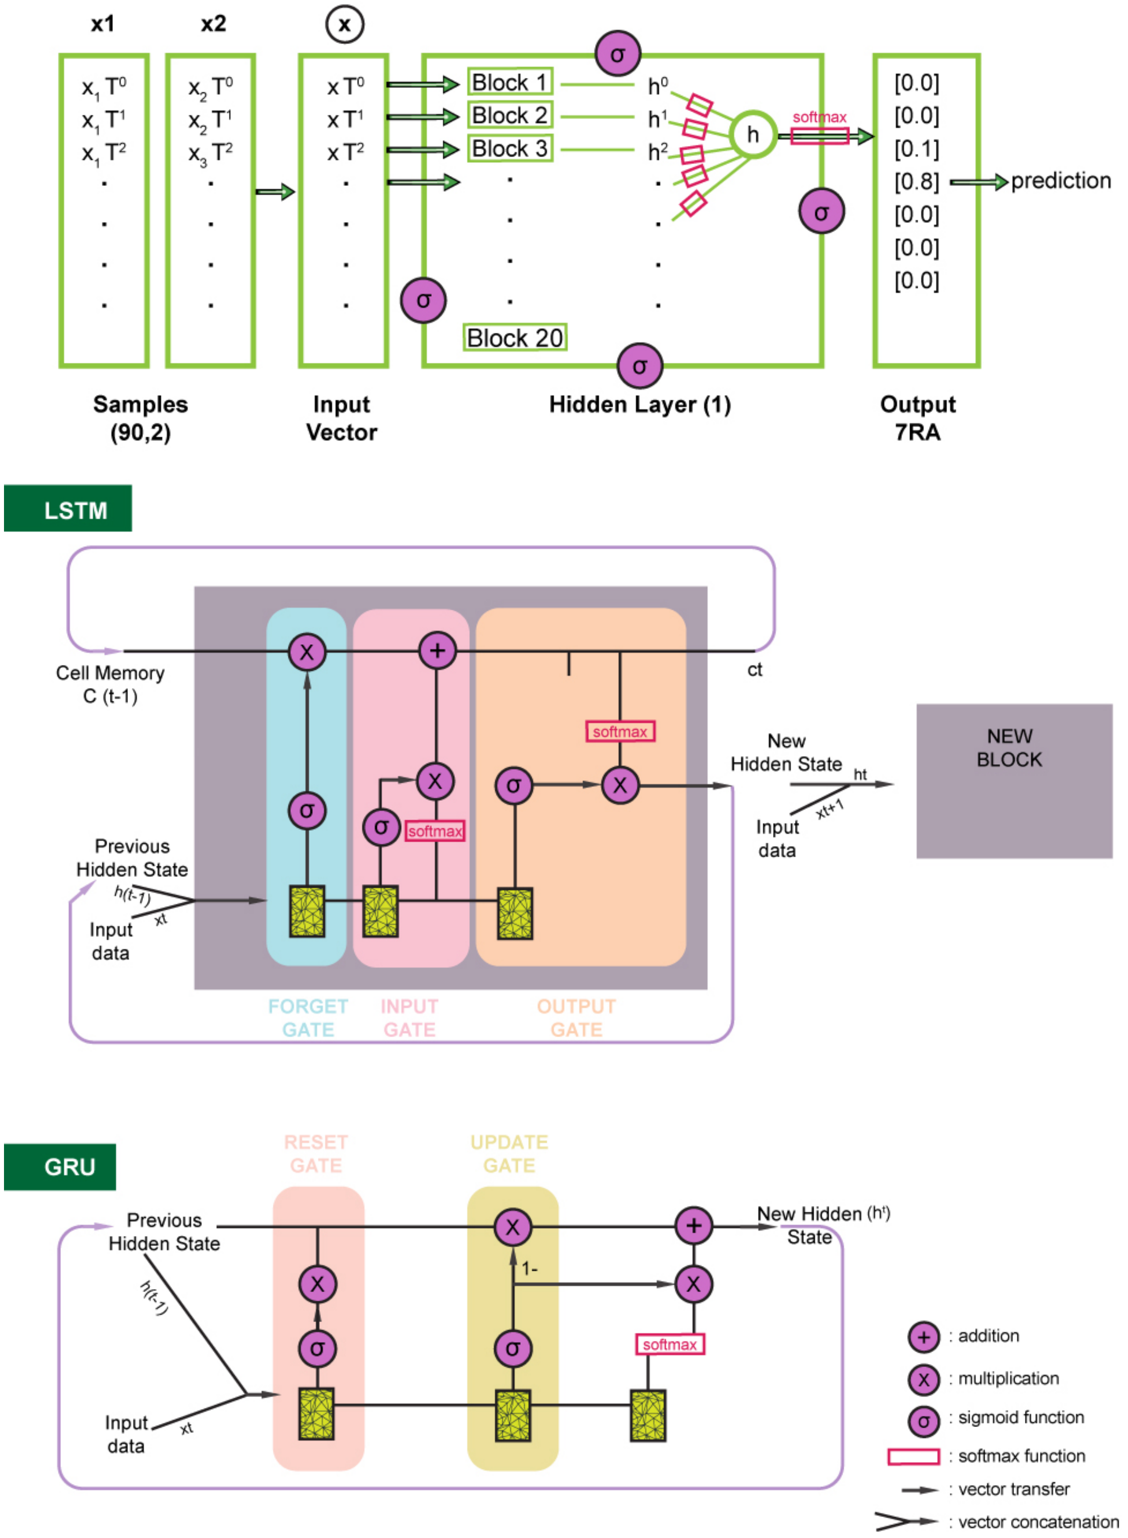

Convolutional Neural Network Model

Our CNN model used a single two-dimensional convolutional layer (Figure 19). CNNs are typically used in image processing but are also popular for classification tasks using sequential data. A layer contains multiple filters that react to different data patterns. The filters are passed over portions of the data, with the user defining the size of those portions and whether the portions can overlap. In our case, the size was 9 by 2, meaning that each portion contained data from both L and R choice responses after nine sequential reinforcers – those reinforcers could be across two different sessions. Following the convolutional layer, it is standard practice to use pooling to aggregate that layer's output and reduce the dimensionality of the data. Here we used max pooling with a pool size of 2 by 1. This means that data were combined column-wise (i.e., max(t0, t1), max(t2, t3) and so on) for each choice. Following pooling, data was converted from a matrix into a single vector (a.k.a. flattened) before being input into the final layer.

The final layer of each model was a densely connected layer with seven nodes, one for each possible label in the original data. The output of this layer was passed through the SoftMax function. The function's output is a final vector with seven elements between 0 and 1. These elements are the probabilities that the model uses to predict the original value of the labelled data, with the highest value being the one predicted by the model. For example, an output of [.05, .05, .1, .1, .2, .3, .2] would mean the model's prediction is 6 as the highest value .3 is the 6th element. During training, these output values were compared to the one-hot encoded outputs to calculate the loss, with sparse categorical cross-entropy used as the loss function. A summary Table 9 that follows, presents the architectural characteristics utilised for training.

Figure 19

Schematic representation of the convolutional neural network with a single hidden layer.

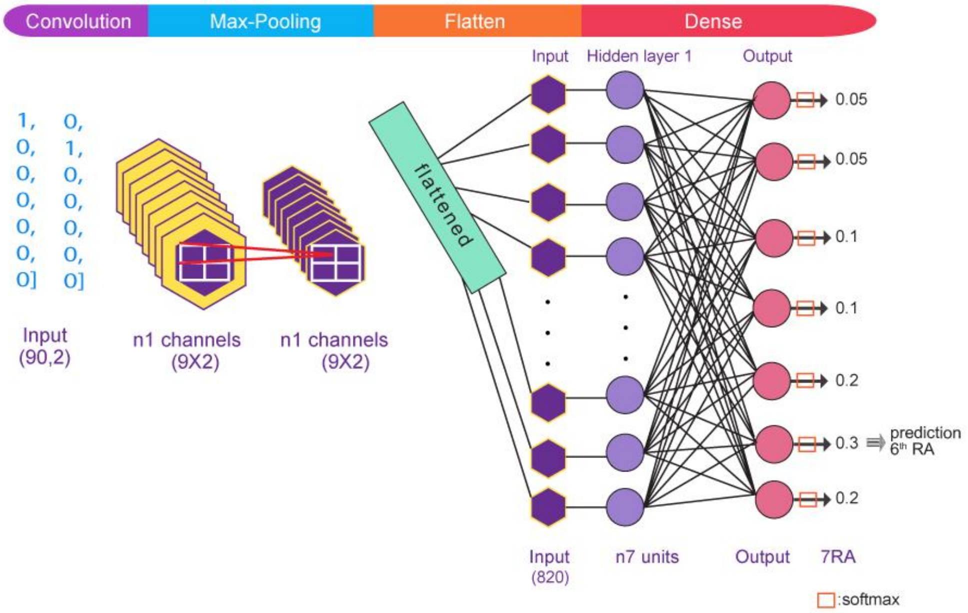

Table 9

Deep Neural Networks Parameters

Type	Output	Settings	Activation	Trained Parameters
	[90, 2]			
LSTM	20	Units = 20	Sigmoid	1,840
Dense	7	Units = 7	Softmax	147
				1,987
	[90, 2]			
GRU	20	Units = 20	Sigmoid	1,440
Dense	7	Units = 7	Softmax	147
				1,587
	[90, 2]			
Conv2D	[82, 20]	Filters = 20, Size = (9,2)	Sigmoid	380
Max Pooling	[41, 20]	Pool Size = (2, 1)	-	-
Flatten	820	-	-	-
Dense	7	Units=7	Softmax	5,747
				6,127

5.2.3 Training procedure for the RNNs and CNN

Unlike other statistical methods, the parameters of a deep neural network are user-defined, meaning an infinite number of models can be tested. To select a final model, techniques such as a grid search can find the best model. A range of variable parameters is chosen in a grid search, and several models are trained using different combinations of those parameters. For the RNN models, the number of LSTM and GRU units varied during the search, as well as the activation function for the layer. For the CNN model, we fixed the convolutional shape and pool size parameters for our experiments. Then, we ran a grid search for all models to find the best-performing model. The "best" was defined as obtaining the highest average accuracy for the test data. The same rule to extract the data was implemented to make results comparable to the single-neuron SNN developed by Plessas et al. (2022). In their study, classification accuracy was calculated following three steps: a) calculated the average firing rate of the training samples belonging to each reinforcement arrangement; b) labelled the predictive reinforcement arrangements by the SNN for each sample with the reinforcement arrangement having the closes firing rate; c) finally compared the predicted and actual reinforcement arrangement of all samples by calculating the errors and then the classification accuracy (Figure 20).

Figure 20

A schematic representation of average firing rates calculation using Euclidian distance where C_1, C_2…C_7 *represent the reinforcement arrangement for prediction.*

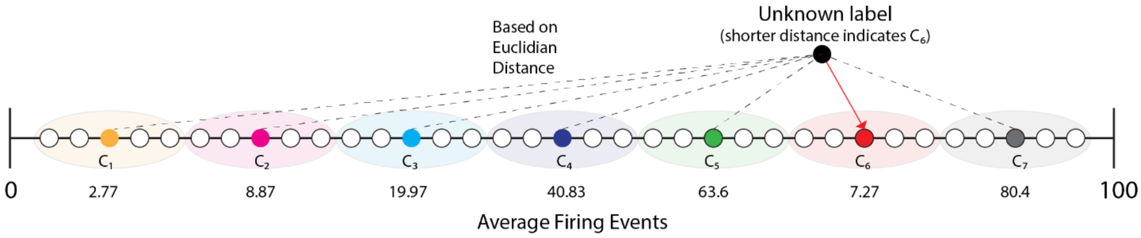

During the search, training the data were divided into a training set and a test set. The training set contained all samples from five pigeons, whereas the test set contained the samples from the remaining pigeon. This process was repeated six times for each search trial so that all pigeons served as the holdout data. Furthermore, each trial was repeated ten times, and an average of the data was taken for selecting the final model. To prevent overfitting to the training data, we stopped training the model once the loss calculated for the test data stopped decreasing for 50 iterations. The model was then reverted to the state that produced the lowest training loss (i.e., 50 iterations earlier).

5.3 Results and discussion

5.3.1 Deep neural networks results

1. Overall performance of the classification models

We used all data from all six pigeons’ datasets for training and testing our models. We selected the six models that showed the highest accuracy in predicting the reinforcement arrangement that shaped the pigeons’ choices.

Table 10 shows that, on average, the combined training and testing accuracy measure in all three deep neural network architectures (CNN, LSTM, & GRU) was between 89–94%. Our models were trained on five pigeons' data and validated on the remaining pigeon that formed the test set. These overall results indicate that DNNs can learn to predict the reinforcement arrangement that was in effect from pigeons’ present binary choices.

Table 10

Accuracy Results of Best Performing Models for all Testing Combinations

Models	Tested on:	CNN	LSTM	GRU
M1	Pigeon 61	0.94	0.91	0.93
M2	Pigeon 62	0.93	0.93	0.93
M3	Pigeon 63	0.90	0.91	0.89
M4	Pigeon 64	0.92	0.92	0.93
M5	Pigeon 65	0.93	0.91	0.93
M6	Pigeon 66	0.94	0.93	0.94

Note: After training the models, the tests were conducted on unseen data. For example, M1 was trained on Pigeon 62-66 datasets and tested on Pigeon 61.

We then compared training and testing outcomes for all models' performances (Table 11, Table 12, Table 13) to identify any overfitting issues for each DNN. When comparing the accuracy performance for all three neural networks (Table 11, Table 12, Table 13), we found that the CNN models overfitted the training data more than the other two RNNs (LSTM, GRU). The differences in accuracy performance between the training and testing results were considerable for CNN models, with training values approaching 100% accuracy while most test values were below 90%. On the other hand, the RNN models had an average training value for accuracy of around 94% and an average test value above 90%, indicating that the ability to predict beyond the training dataset (i.e., generalisation ability on novel data) was higher for LSTM and GRU-based models. It was then vital to determine whether those results were due to the specific architectures or other factors such as hyperparameter tuning. Using different hyperparameters could have led to less overfitting in the CNN-based models or more overfitting in the RNN-based models. For example, Table 9 shows that the CNN had at least three times as many trainable parameters as the RNN-based models. Increasing parameters can often lead to overfitting; therefore, although all DNN models reached high accuracy, the CNN-based method is overfitting, and its generalisation ability decayed.

Table 11

Results of Best Performing CNN Models for all Combinations when Splitting the Data

	M1	Pigeon 61	M2	Pigeon 62	M3	Pigeon 63	M4	Pigeon 64	M5	Pigeon 65	M6	Pigeon 66
	Train	Test	Train	Test	Train	Test	Train	Test	Train	Test	Train	Test
Accuracy	0.99	0.89	0.98	0.88	0.96	0.85	0.99	0.85	0.97	0.90	0.99	0.89
Recall	0.96	0.62	0.94	0.60	0.87	0.48	0.98	0.50	0.91	0.65	0.96	0.63
Specificity	0.99	0.93	0.99	0.93	0.97	0.91	0.99	0.91	0.98	0.94	0.99	0.93
Precision	0.97	0.65	0.94	0.63	0.87	0.49	0.98	0.49	0.92	0.72	0.96	0.67
F1	0.96	0.60	0.94	0.58	0.87	0.46	0.98	0.56	0.91	0.62	0.96	0.63
Informedness	0.96	0.55	0.93	0.53	0.85	0.38	0.97	0.41	0.90	0.59	0.96	0.57

Table 12

Results of Best Performing LSTM Models for all Combinations when Splitting the Data

	M1	Pigeon 61	M2	Pigeon 62	M3	Pigeon 63	M4	Pigeon 64	M5	Pigeon 65	M6	Pigeon 66
	Train	Test	Train	Test	Train	Test	Train	Test	Train	Test	Train	Test
Accuracy	0.91	0.92	0.93	0.93	0.92	0.91	0.93	0.92	0.93	0.90	0.94	0.93
Recall	0.70	0.74	0.78	0.76	0.72	0.69	0.77	0.72	0.78	0.66	0.81	0.75
Specificity	0.95	0.95	0.96	0.96	0.95	0.95	0.96	0.95	0.96	0.94	0.96	0.95
Precision	0.71	0.72	0.78	0.79	0.74	0.71	0.78	0.75	0.78	0.66	0.81	0.79
F1	0.69	0.79	0.77	0.77	0.72	0.70	0.76	0.71	0.77	0.69	0.80	0.74
Informedness	0.65	0.70	0.75	0.72	0.68	0.64	0.73	0.68	0.74	0.60	0.78	0.71

Table 13

Results of Best Performing GRU Models for all Combinations when Splitting the Data

	M1	Pigeon 61	M2	Pigeon 62	M3	Pigeon 63	M4	Pigeon 64	M5	Pigeon 65	M6	Pigeon 66
	Train	Test	Train	Test	Train	Test	Train	Test	Train	Test	Train	Test
Accuracy	0.94	0.93	0.94	0.93	0.93	0.86	0.94	0.92	0.94	0.89	0.94	0.92
Recall	0.79	0.78	0.81	0.76	0.77	0.53	0.82	0.72	0.79	0.62	0.81	0.73
Specificity	0.96	0.96	0.97	0.96	0.96	0.92	0.97	0.95	0.96	0.93	0.96	0.95
Precision	0.80	0.81	0.81	0.78	0.77	0.52	0.82	0.74	0.79	0.60	0.81	0.80
F1	0.79	0.77	0.81	0.77	0.76	0.60	0.81	0.71	0.78	0.65	0.80	0.73
Informedness	0.75	0.74	0.78	0.71	0.73	0.44	0.79	0.68	0.76	0.56	0.78	0.69

Lastly, we calculated additional performance measures, keeping in mind that accuracy performance omits other data characteristics (Kuhn & Johnson, 2018) that can be crucial in decision-making processes. As Plessas et al. (2022) have discussed, choices related to specific reinforcement arrangements should create unique patterns that differentiate them from choices shaped in other, different training environments. By considering other performance metrics, we evaluated the ability of each model to identify true positives and negatives. In particular, when classifying binary choices, for the outcome to have significance in behavioural research, the model needs to discriminate between events that belong to the same reinforcement arrangements (choice related to the environment) and events that do not (i.e., choices that are not related to the environment).

While considering additional measures of performance, we noticed that the CNN-based method failed two tests when using data generated by Pigeons 63 and 64: measures of recall dropped significantly (48-62%). This outcome implies that overfitting was occurring, as mentioned above. When we compared LSTM and GRU-based models, GRU did worse in predicting the reinforcement arrangement in effect (Table 13). Though accuracy measures show that LSTM and GRU (Table 12) can be trained by binary choices to identify the reinforcement history, the variability in other performance measures may indicate that some optimisation is needed for sequential prediction problems with binary choices as inputs. Lastly, we noted that specificity was higher than recall, indicating that all models did better in identifying choices that had very distinct behaviours (i.e., when behaviours were less alike); therefore, RNNs were better at determining what training environment the choice behaviour did not belong to rather than which environment it did belong to.

2. Accuracy performance per reinforcement arrangement

By the following analysis we investigated how accurately each model predicted each reinforcement arrangement in effect for models M1–M6. Again, the same general trend was seen as in the previous analysis. Figure 21 shows the models' performances for each pigeon when multiple trials were used for the testing data. The groupings of data points suggest that while there was some variability between model runs, the models could typically identify each reinforcement arrangement with at least 80% accuracy for each dataset used for validation. The only exception to this was Reinforcement arrangement 3 (for Pigeon 64, all runs of the CNN-based models were less than 80% accurate at predicting that particular reinforcement arrangement). It is unclear why this was the case, as that reinforcement arrangement was accurately identified by the other model types using Pigeon 64 datasets as the validation test.

Figure 21

Accuracy for all seven reinforcement arrangements for each model (M1-M6) when tested on Pigeons 61–66 accordingly.

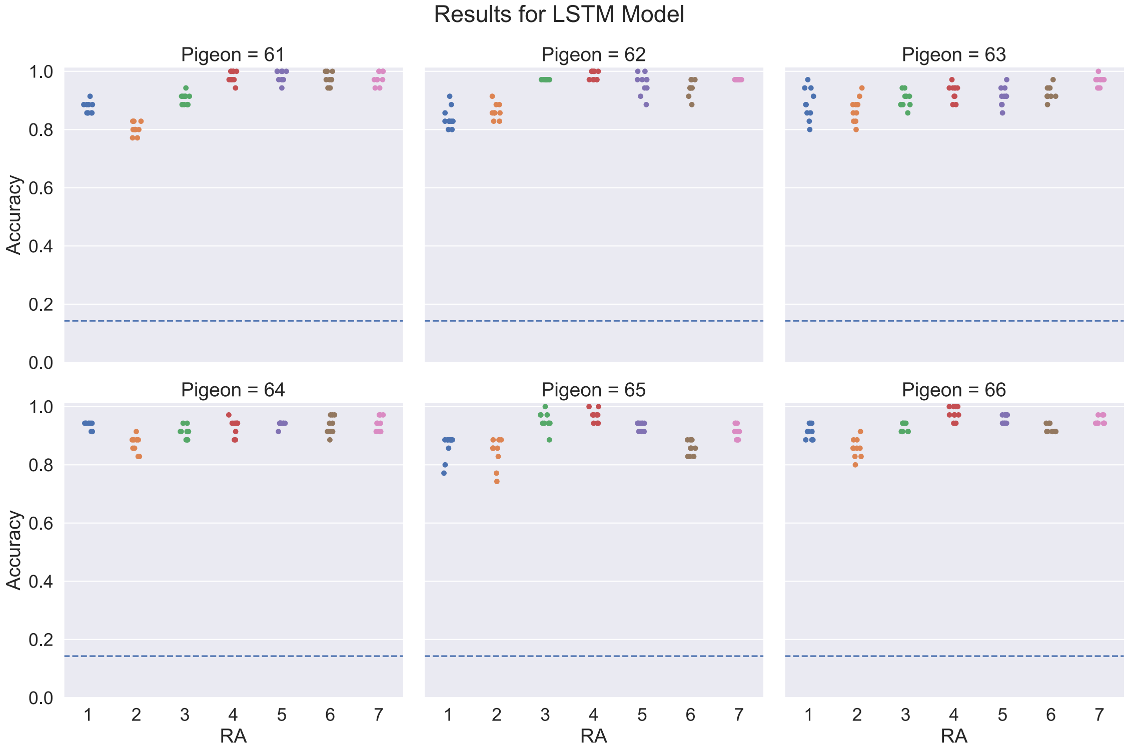

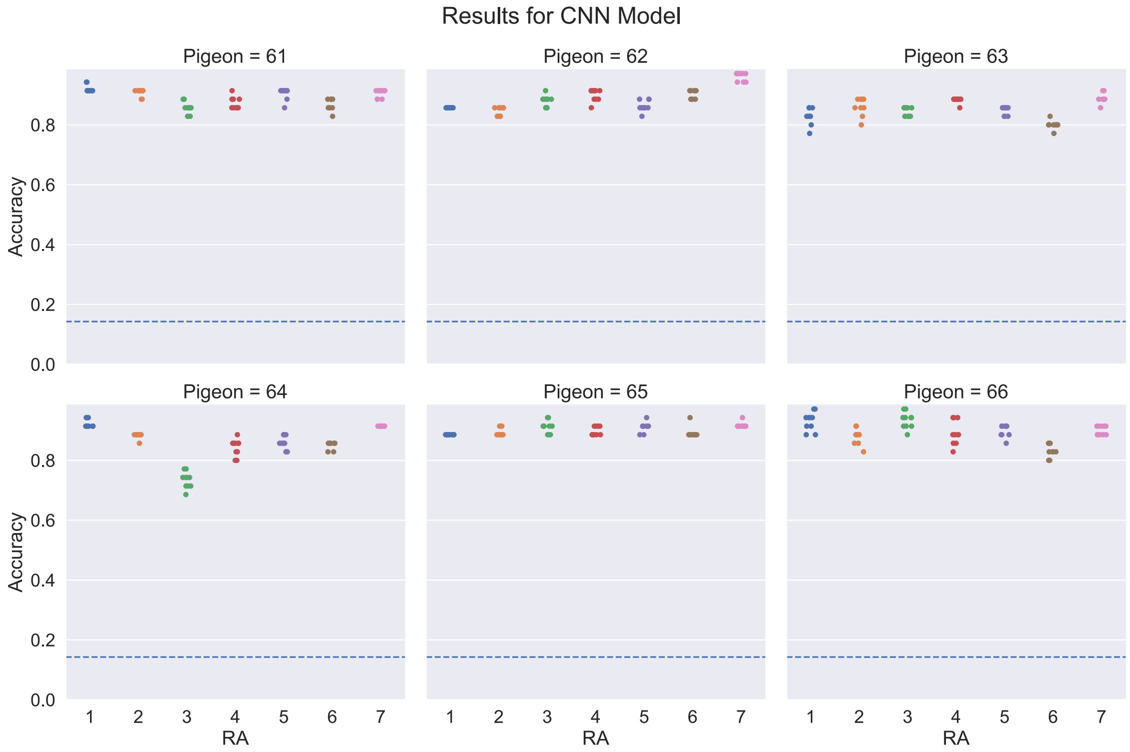

Note: RA denotes the seven reinforcement arrangements

We produced confusion matrices for Pigeons 63 and 66, as these two subjects tended to produce the lowest and highest validation accuracy, respectively. The models had to be trained to deliver the matrices because the required information for the matrices was not collected during the original experiment. Figure 22 displays the confusion matrices as heatmaps. These figures show that when an error was made in classifying a reinforcement arrangement, the actual reinforcement arrangement was confused with another, closely related reinforcement arrangement. In other words, data from Reinforcement arrangement 2 could be confused with Reinforcement arrangements 1 or 3, but never with more distant reinforcement arrangements.

Figure 22

Confusion matrices for Pigeon 63 and Pigeon 64 that produced the lowest and highest validation of accuracy, respectively.

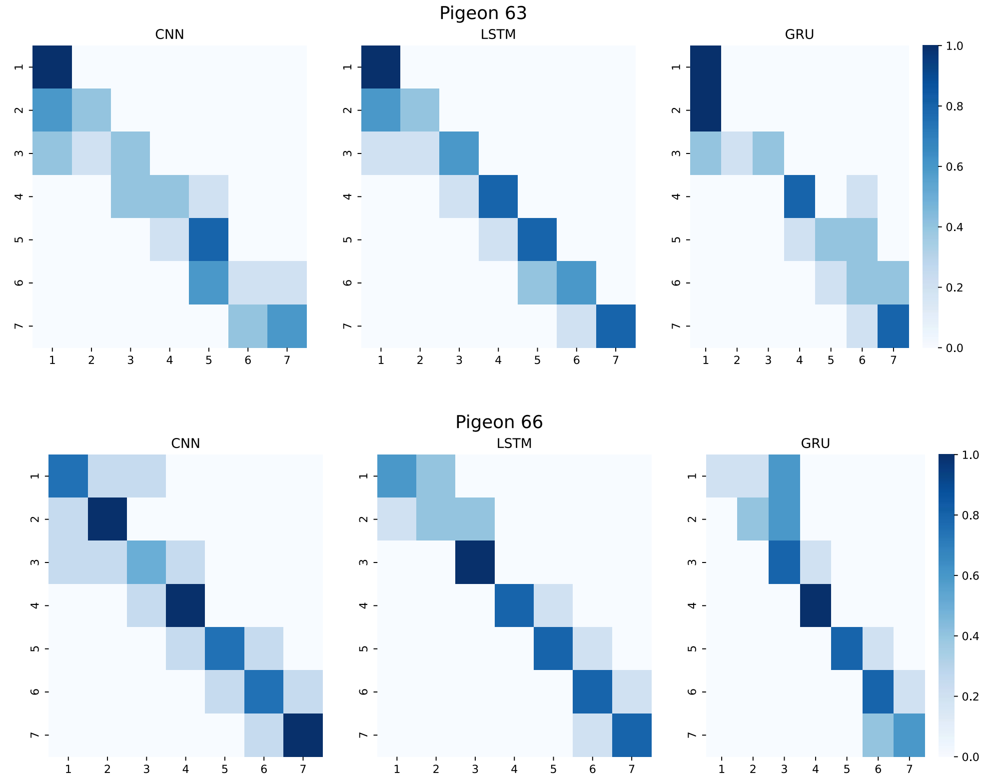

Note: The vertical line in each graph denotes the predicted values of the seven reinforcement arrangements and the horizontal, the actual values.

3. Comparing the results to SNN

We compared the performances of all the DNN models that we trained (CNN, LSTM, GRU) with the performance of SNN as published in Plessas et al. (2022) by using accuracy and F1 scores as a weighted average of precision and sensitivity. Though the most common measure is accuracy, we decided to include F1 to account for additional data characteristics captured by sensitivity and specificity metrics. Figure 23 demonstrates that when looking into the average accuracy (training and testing data included), all models could predict the reinforcement arrangement that shaped the pigeon's behaviour with an average accuracy above 88%.

Figure 23

Accuracy score performances for all neural networks (best models).

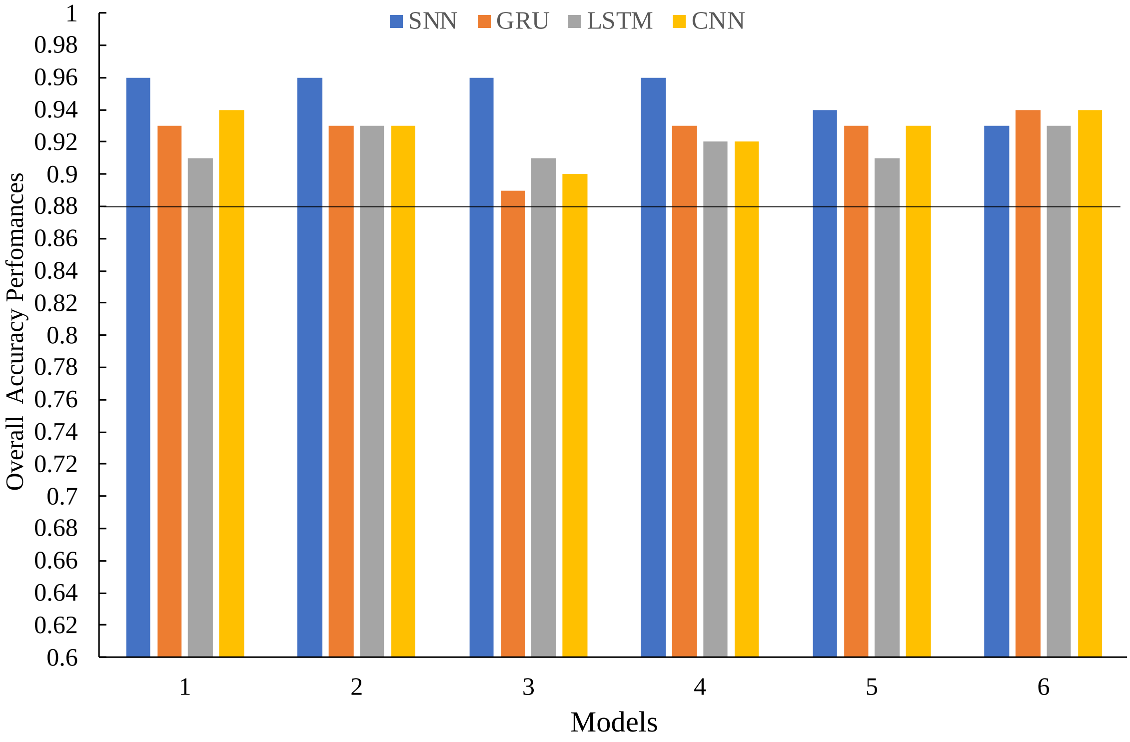

When the F1 measure is considered, the same outcome was not demonstrated (Figure 24): SNN outperformed the prediction ability of all other trained models. All DNN models (CNN, LSTM, GRU) performed between 66–79%, indicating that all models might have benefited from optimisation to achieve better prediction in discriminating events that look alike and could share similar patterns. At the same time, no further training was required with SNN, and it could successfully handle the data.

Figure 24

F1 score performances for all neural networks (best models).

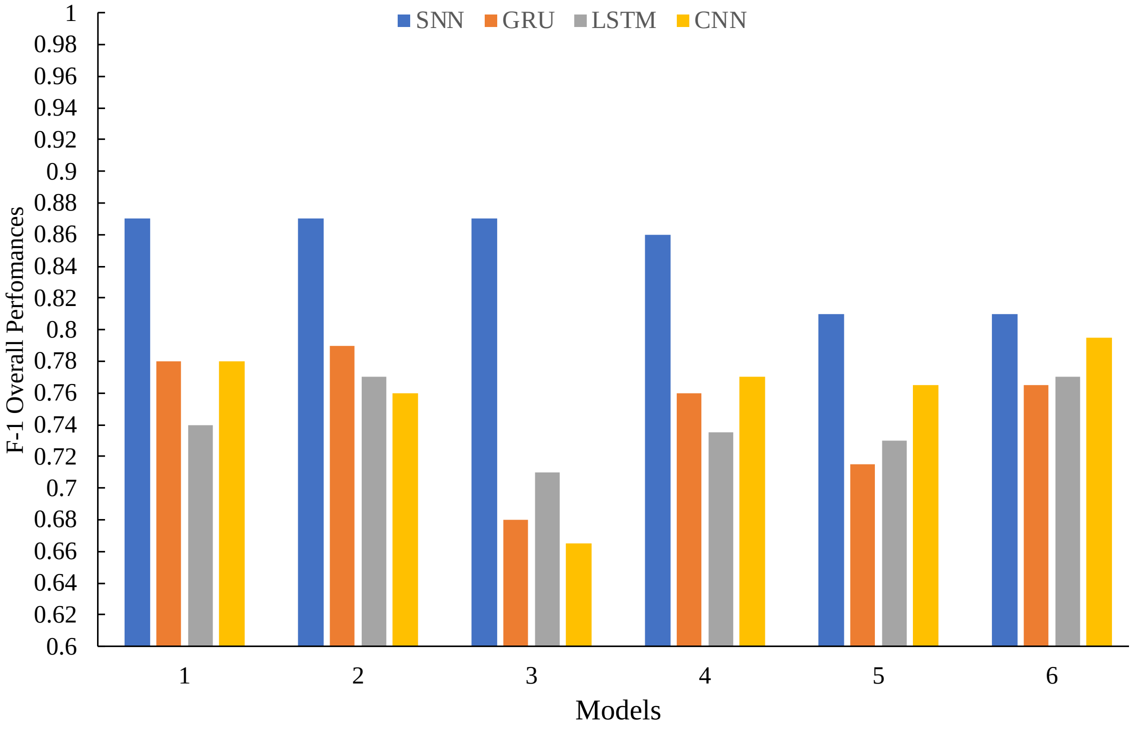

5.4 Overall discussion

Behavioural research has demonstrated that different reinforcement arrangements maintain different patterns and rates of behaviour, and reinforcers can be manipulated to influence current behaviour. In this study, we investigated whether a reverse test was sufficient when we used pigeons' binary choices as input data for deep learning neural networks to generate predictions of the reinforcement arrangements in the behavioural experiments that caused a particular behaviour that occurs in current time. Our study demonstrated that deep learning neural networks could handle binary choices as temporal behavioural research variables and detect patterns in current behaviour related to past reinforcement arrangements. Furthermore, this relationship between behaviour and environment was detectable in datasets containing a minimal amount of data (5-sec). This outcome strengthens what behavioural research has long demonstrated, i.e., that behaviour is influenced by the environment in effect, and validates the theory that learnt behaviour carries information (hidden detectable patterns) that informs us about this relationship. Approaching behavioural research data by conducting reverse tests as in this study could enrich our understanding choice behaviour. Also, it shows how reverse tests work with different types of neural networks allowing us to investigate their performance and generalise ability with real datasets and helping to guide the development of more effective neural networks in the future.

Not every model performed equally well when it was evaluated on unseen data. CNN-based models showed an overfitting effect with an approximate 10% drop in accuracy on the tested datasets. Also, when we included additional measures of prediction performance, two CNN-based models failed to predict the reinforcement arrangement. As discussed by (Plessas et al., 2022), a model needs to perform beyond accuracy and do well in specificity and recall measures of performance to conform with the generalised matching law (Baum, 1974). Thus, as variability in data is idiosyncratic, it becomes necessary to look into all performance metrics, not just accuracy. Sequential processing with recurrent models (LSTM, GRU) was more successful in all measures than CNN. This accords with the fact that RNN-based models are more capable of interpreting temporal data than CNN-based models, as their architecture deals with temporal variables via sequence processing. As both layer types contain gates that determine how much information is passed to the following part of the sequence, future research may look into how the amount of information that passes through the gate may affect the ability to identify the reinforcement arrangement that was in effect.

Further, when looking into the ability of each model to predict, based on the reinforcement arrangement used, the analysis revealed that all artificial neural networks were better at identifying patterns trained in more distinct environments (Reinforcement arrangement 1 and 7), whereas confusion occurred with less disparate reinforcement arrangements which were programmed to have more similarities.

When the trained models were compared to SNN performance, we observed that the SNN models developed in (Plessas et al., 2022) outperformed all models trained in this study (CNN, LSTM GRU), indicating that SNN may be a better fit when binary choices are used as inputs. For DNNs to reach similar outcomes, it may be necessary to use optimisation strategies. SNN did not require additional training, as optimisation is a built-in process during training, excluding the need for additional steps to reach acceptable predictive outcomes with unseen data. With SNN, the authors had addressed the training as an optimisation process, where the objective function was to find the vector of weights that maximised the classification accuracy of the input patterns. The use of a differential algorithm in training sufficed for the model to make decisions on current behaviour-environment interactions.

Moreover, SNN did not require any pre-processing technique for the data to be used as inputs, whereas pre-processing the data was necessary with DNNs; SNN treats data as discrete events that occur at certain points in time. The pre-processing steps needed for other types of neural network slightly modify the data, which in turn affects the predictive performance of each model we created. Depending on the method used, adding a step to pre-process the data could affect the training procedure and the output could miss some of the dimensionality of the data (i.e., spatiotemporal values). Future research can investigate how different pre-processing strategies affect sequential processing. This may play a role not only with binary choices but even with choices from more than two alternatives.

A limitation of this study is the different training strategies we used to develop the neural networks; with SNN modelling, we used cross-validation, while with DNNs, a holdout procedure. It is possible that if the same training procedure were used for each architecture, the outcomes would be different as each algorithm responds differently to training procedures, and there is no way to predict if, by using similar training procedures, we would have seen the best performance for each algorithm. The decision as to what training procedure to use was based on common practice in the field for training each type of neural network.

To summarise, all three types of neural networks were able to produce positive outcomes, though depending on the architecture of the neural network, the training procedure, and the input data, their predictive ability differed. Having this in mind, one should consider that with behavioural data, the selected algorithms should produce an outcome that fits with clinical decisions. Extensive research has demonstrated binary choices are dependent on the recent reinforcement history as well as other factors such as the location in time of the behaviour relative to the reinforcer delivery and the overall rate of reinforcement (Alsop & Elliffe, 1988; Davison & Baum, 2000). Therefore, temporality is a crucial variable better handled from SNN and RNN-based models. It would be interesting to investigate which model would maintain the highest predictive ability, for example, if more choices were given as an option (e.g., three, four, five, six) or if more datapoints in time were used (i.e., beyond the 5-seconds rule set in this study). We assume that varying the amount of choices used from different points in time and location would add a more complex spatial element to our data, which, theoretically, would favour the use of SNN that is designed to handle spatiotemporal variables as discrete events rather than treating the data exclusively on the sequence they occurred (i.e., RNN-methods). As the use of such data with AI research is still new, further experimentation will be needed to determine the parameters for best predictions in clinical research with choice behaviour.

5.5 Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

5.6 Data availability

The datasets used in this study were generated by Plessas et al. (2022) which authors made available upon request. The experimental code for the main experiment can be found following the link: https://anonymous.4open.science/r/behavioural-neural-networks-3820

5.7 Links between Chapters 5 and 6

Overall, different neural network architectures differ in their predictive ability. The paper presented in Chapter 5 demonstrated that the single-neuron SNN developed in Chapter 4 outperformed the DNN models, indicating that the event-based processing used with SNN as it stands is a more useful approach for analysing choice behaviour data in concurrent schedule procedures. This outcome by no means implies that the single-neuron model was the best rather, that SNN could handle the data successfully without needing any further training and thus validating the approach we had taken so far. At the same time, DNNs might improve through some sort of optimisation process, which was not the focus of this thesis, as the goal was not to find the best model but rather a useful model to solve the problem of interest.

Hence, the next step required further investigation with other datasets. A useful single-neuron SNN would be able to generalise the outcome, to predict the reinforcement history by conducting a reverse test to other populations, groups or individuals, even though their choices were learnt in similar –but not identical– environments. By including datasets from other studies, external validation of the SNN model would be strengthened, but also the predictive outcome would allow us to investigate further environment-behaviour interactions. Table 2 on page 64 provides a summary of contributions.

Chapter 6 A spiking neural network approach to personalise learning histories for pigeons' choice-making

Prelude

Behaviour analysts are primarily interested in studying adaptive learning by analysing response patterns unique to each individual. Through carefully controlled experimental designs, they rely on strictly observable phenomena to investigate the impact of a given environmental condition on behaviour. Small numbers of subjects are used (small-N research) usually using within-subject designs with a strong emphasis on experimental design and replication. Therefore, the generality of results is achieved through replication (Sidman, 1960) (for more, see 2.6).

Therefore, any AI tool that is developed to predict learning histories from choice behaviour data cannot be approached as a one-size-fits-all technology. In the pilot study presented in this chapter, we are extending the research findings from Chapter 4 to investigate if an SNN model can generalise to new datasets drawn from a different study, one with different animal subjects, and procedural differences, than the ones used to create the artificial model. In Chapter 6, the focus was to personalise the model outputs by considering the unique patterns represented in each organism's choice behaviour dataset (Figure 25).

Figure 25

A visual representation of how new datasets were used to evaluate the SNN performance.

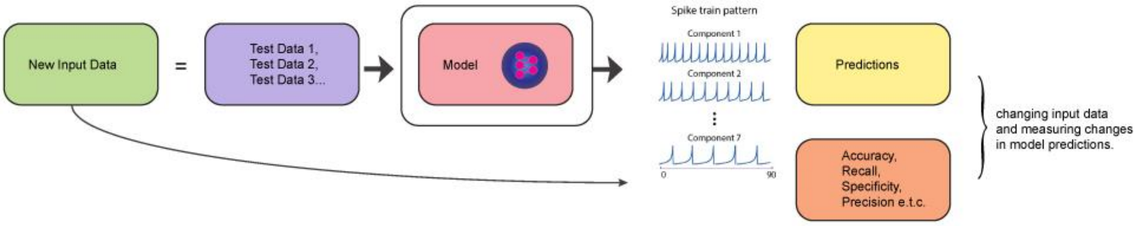

At the beginning of this chapter, there is an opening section introducing the additional datasets used for generalisation. This section outlines the differences between the procedure used in the additional source experiments (Davison & Baum, 2000) to the dataset used in Chapter 4. This information is beyond that included in the manuscript presented in this chapter.

Pigeon choice behaviour datasets from Davison and Baum (2000) formed the basis for the study presented in this Chapter. Permission was obtained by the authors to use these datasets for this project (Appendix A). Some additional experimentation was done with other species, rats and humans in testing choice behaviour datasets. These data were shared with the first author's permission. As the data extracted were only used for preliminary testing, they were included in (Appendix C) and will not be described in this section.

Additional datasets

New datasets other than the ones used in Chapter 4 were included in the experiment in this chapter. The datasets were extracted from Davison and Baum (2000), where they manipulated the training environment to measure six pigeons' choice making. Their research adapted a published procedure (Belke & Heyman, 1994) where the training environment changed rapidly to measure how this process affected choice behaviour. As in the study described in Chapter 4 (Prelude), Davison and Baum (2000) arranged seven reinforcer ratios used in a two-key concurrent VI schedule. Each VI-VI pair was presented in random order and separated by 10-sec blackouts. The difference in their procedure to the source study was that they varied the number of reinforcers given for each component between four and twelve, whereas Landon and Davison (2001) kept that number constant (ten reinforcers). Thus, the speed of environmental change varied. In addition, they altered the overall reinforcer rate within components, from four to eight reinforcers, whereas Landon and Davison (2001) kept the overall reinforcer rate constant (Table 14).

These procedural manipulations made the change in the environment faster in components and affected the sensitivity the choice behaviour had towards reinforcement (for more, see 1.2.5). The rate of choice behaviour change, in such a rapidly changing environment, was higher when the overall reinforcer rate was higher. They also noted that most changes in choice behaviour were at the beginning of their exposure to new environmental conditions rather than later. Therefore, choice behaviour changed initially quickly in response to the delivery of reinforcers and then stabilised, leading the authors to conclude that events more distant in the past may have less impact on current behaviour. At the same time, the number of reinforcers per component had little discernible effect.

Table 14

The Differences and Similarities in the Source Experiments Environmental Training Procedure

Datasets extracted from	Participants	Reinforcers per un-signalled component	Reinforcers per minute
Landon & Davison (2001) Experimental Conditions: C1, C6	Pigeon 61 to 66	10 [maximum 70 per session]	2.22
Davison & Baum (2000) Experimental Condition: C1, C4, C 9	Pigeon 91 to 96	10 (C1) and 4 (C4, C9) [maximum 70 (C1) and 56 per session (C4, C9)]	2.22 (C1, C4) and 6 (C9)

Manuscript

Plessas A, Espinosa-Ramos J, Landon J, Cowie S, Parry D. A spiking neural network approach to personalise learning histories for pigeons' choice-making. To be submitted June 2023

6.1 Introduction

Learning is demonstrated in the outcomes of the dynamic relationship between organism and environment, as the environment shapes the organism's interactions with it. Behaviour analysts have extensively studied the functional relations between environment and behaviour by manipulating reinforcer arrangements and observing changes in the behaviour of organisms (Cowie & Davison, 2016a). The goal of this study was to evaluate the ability of a machine learning (ML) tool, trained on a behavioural dataset of binary choices, to generalise to new datasets. Specifically, we sought to determine whether the tool could retrodict learning histories from individual response patterns that occurred within the first five seconds after the delivery of a reinforcer.

In behavioural experiments, one way to investigate how a behaviour changes is to focus on choice behaviour in concurrent schedules that change frequently and variably to simulate the natural environment. In concurrent-schedule procedures, the participant distributes its response or time between two available alternatives in accordance with the distribution of reinforcers obtained in these two alternatives (Baum, 1974, 2010; Herrnstein, 1970). Local choice analyses describe how the allocation of behaviour changes as a function of time or responses since the most recent reinforcer. Generally, choice behaviour favours the alternative more likely to deliver a reinforcer at any given moment in time. In a rapidly changing procedure, reinforcers tend to induce a period of extreme preference toward the response that produced them (e.g., Davison & Baum, 2000, 2002). In more stable environments, local choice favours the response that produced it only when that response is also more likely to produce the subsequent reinforcer (e.g., Cowie et al., 2011; Krägeloh et al., 2005). Although some component of this analysis is an artefact of the reinforcer interrupting a visit (McLean et al., 2014), these preference pulses (i.e., the strong tendency for the first few seconds toward the same alternative) do reflect an actual effect of reinforcers on responding only when future reinforcer availability is signalled (Gomes-Ng et al., 2017).

In this study, the retrodiction of an artificial model is explicitly based on a fraction of individualised choice behaviours (5-second periods after the delivery of each reinforcer) from experiments where the reinforcers had different functions. A useful artificial model that learns from these datasets should generate variable retrodictions that capture differences in response patterns affected by reinforcer effects over time. The approach is novel and significant both from a theoretical perspective and because it has implications for practice. From a theoretical perspective, the approach holds great promise in advancing the study of single-organism behaviour as it allows for an exploration of the connections between learning histories and subsequent reinforcers acting as markers in time. Also, it could offer valuable insights into the intricate dynamics of how learning experiences shape behaviour in individuals. Therefore, it can be used as a tool to understand both generality across species, and individual differences. Furthermore, in clinical practice, retrodiction can point to otherwise unknown learning histories of individual-specific influences on the extent to which learning takes place. Thus, retrodiction can inform the development of future specific optimal environments for learning by identifying learning histories. As Staddon (2016) points out, group data cannot be extrapolated to all individuals within a group or between different groups.

In Plessas et al. (2022) unknown learning history was able to be identified by a brain-inspired, spiking neural network (SNN) based on current binary choices. This finding confirmed the well-established relation between choice behaviour as it evolves over time, and reinforcers that are contingent upon these choices. The SNN, which consisted of a single Leaky Integrate-and-Fire neuron, produced accurate retrodictions using a snapshot of pigeons' binary choices in the 5-sec period of choices made after reinforcer delivery, from two experimental conditions extracted from Landon and Davison's (2001) study. The SNN was trained and validated using one dataset (Condition 1), and its performance was tested on a separate dataset (Condition 6) to evaluate its ability to generalise to new situations(Plessas et al., 2022).

Extending this approach, the purpose in this study was twofold: First, to evaluate the performance of the single-neuron SNN model on new datasets obtained from a different study (Davison & Baum, 2000). The datasets extracted, consisted of choice behaviours related to identical and non-identical environments to the ones on which the single-neuron SNN model was developed. Second, to investigate if the single-neuron SNN model can be personalised by identifying the unique patterns developed by each organism during learning. *Personalised modelling* has many functions in machine learning (Fan & Poole, 2006). In this study, personalised modelling is a focused generalised task. The retrodictions of the SNN aims to address individual variability (focused) while considering insights into the common aspects of inter-individual variability (generalised) in the neuron's spiking activity. By incorporating both perspectives, the study aims to provide a comprehensive understanding of the neuron's spiking activity at both the individual and group levels.

Therefore, the firing rates produced by the spiking neuron differ for each individual (and for a specific group). The firing rates observed in the new datasets also differ from those generated when participants made choices in the study used to develop the single-neuron SNN model (Plessas et al., 2022). As a result, the novelty of this approach lies in accounting for personalised learning histories by using new firing rates, as opposed to the ones that were used to train the model. Thus, in part the ‘behaviour’ of the spiking neuron changes as the frequency of the spikes generated by the neuron are adjusted to individual variability.

6.1.1 Related work to personalised modelling

Overall, machine learning (ML) algorithms are trained to identify relationships between inputs and outputs, relationships that exist in the form of patterns. ML analysis of experimental data primarily focuses on prediction, and it has been proposed that it can complement traditional statistical inference techniques (Orrù et al., 2020). However, how ML can complement the analyses of experimental data from single-subject research techniques has received little attention. In most cases, traditional ML modelling relies on large training datasets (usually annotated datasets) where training and testing of the model's performance come from the same dataset. This practice may not apply well to ‘real-world contexts’, or data different from those a model was trained on, as training and testing usually have distinctive qualities (Niu et al., 2021). Training on the same datasets can also create data dependency, the training of a model to become reliant on specific patterns presented in the training datasets, thus, making it less adaptable to unseen or different data. This is one of the most serious problems reported in ML, particularly with deep learning (Tan et al., 2018).

Ribeiro et al. (2016b) recommends that when evaluating the performance of an ML model, individual differences should be considered in analysing the data, as a body of individual predictions could increase our understanding of how the model ‘behaves’. Thus, end-users can make better decisions, trust the model's predictions more, and use them more effectively in their work (Ribeiro et al., 2016a, 2016b). Personalisation with behavioural observations is required because individual behaviour varies in response to events (e.g., in relation to levels of reinforcer discriminability in time), and not all algorithms can capture these differences. A personalised approach takes account of individual differences, performs better at an individual level, and allows for a model tailored specifically for an individual, while the model still learns from available data (Collin et al., 2022; S. Taylor et al., 2017, 2020).

Moreover, (Schneider & Vlachos, 2021) emphasise the significance of understanding both the target and the recipient of model personalisation, as well as the various approaches that can be employed to provide personalised explanations. Different ML techniques have been adopted to achieve personalised predictions and to optimise the performance of the artificial neural network, either for an individual or for a specific group (Schneider & Vlachos, 2021). Transfer learning is an ML method that uses a pre-trained model as a starting point to develop a new artificial model to improve the prediction process (Weiss et al., 2016). Thus, knowledge gained from the pre-trained baseline model is transferred to a target model to improve its ability to make new predictions. A significant advantage of this approach is the reduction of data dependence of the model, and the minimisation of the time the model needs for training, as there is no need to begin training again (Tan et al., 2018). Although transfer learning has been successfully used in traditional deep learning artificial neural networks over the last decade or so, research on transfer learning in SNNs is still lacking (Wu et al., 2023; Yan et al., 2021; Zhan et al., 2021, 2022).

6.1.2 Spiking Neural Networks for personalised modelling

Few attempts have implemented a personalised approach using artificial SNNs (e.g., Breen et al., 2016; Doborjeh et al., 2019, 2022; Saeedinia et al., 2021). The personalised approach in those studies was achieved by training each SNN model from scratch, on a subset of one-participant's static and temporal data, to capture the space and time association patterns. In this study personalised modelling was approached differently by assessing whether the retrodiction can be accurate without additional training by allowing the firing rates to adapt to the new dataset.

In Plessas et al. (2022), a spiking neural network (SNN) was chosen due to its ability to process discrete events that occur in time, by using spikes rather than treating the data as a cumulative number of responses over a continuous period. Thus, the advantage of the algorithm is that it is inherently suitable for temporal data (Yamazaki et al., 2022). Furthermore, SNNs also allow for the spatiotemporal analysis of the data considering the behaviour allocation (left-key versus right-key) from one time to the other, regardless of how the relationship between behaviour and reinforcer was formed and does so by developing a nonlinear model (Izhikevich, 2004). We anticipated the single-neuron SNN to capture differences in response patterns in their firing rates without the network being re-trained. If the model can accurately identify learning histories from new datasets, this would demonstrate the adaptive nature of the 'flexible' neuron and its ability to accommodate new choice behaviours.

6.2 Method

6.2.1 Dataset for single-neuron SNN personalised modelling

The datasets were extracted from Davison and Baum (2000) in which pigeons were trained to make choices in daily sessions composed of seven concurrent-VI schedule reinforcer ratios called *components*. Their procedure showed that participants could change their choice behaviour in relation to the arranged reinforcer regime and, depending on the component, the rate of change in behaviour varied.

We extracted raw data from three experimental conditions (Conditions 1, 4 and 9) for six pigeons, using data from the last 35 training sessions. Each session was organised by setting seven different components with (L[left]: Right[right]) reinforcer ratios as 27:1, 9:1, 3:1, 1:1, 1:3, 1:9, 1:27. We applied the same extraction rule as in the published paper in Chapter 4 (Plessas et al., 2022) to allow comparison of the one-size-fits-all approach/traditional approach, versus personalised approach. Thus, the choice behaviour of the pigeon was to peck to the left (L) or to the right (R) key and was recorded continuously for 5 seconds, as a frequency event after the reinforcer delivery and represented as discrete events in time to feed the artificial model.

In Condition 1, the pigeons received training under the same environmental conditions in each component (n=9 reinforcers, 2.22 reinforcers/min) as in Plessas et al. (2022). In Condition 4, the number of reinforcers delivered per component decreased to approximately half (n=4). In Condition 9, the number of reinforcers was as in Condition 4, but the reinforcer rate was higher (6 reinforcers/min). Therefore, the new datasets were extracted from different subjects (same species) when their choice behaviours were coupled with identical and non-identical environmental training conditions compared to the dataset used for developing the single-neuron SNN model. Further, inclusion of conditions across which the rate and/or the number of reinforcers delivered differed permitted testing the personalised model's performance on datasets formulated by various environmental parameters (Figure 26).

Figure 26

Schematic representation of the choice behaviour datasets used for personalised modelling.

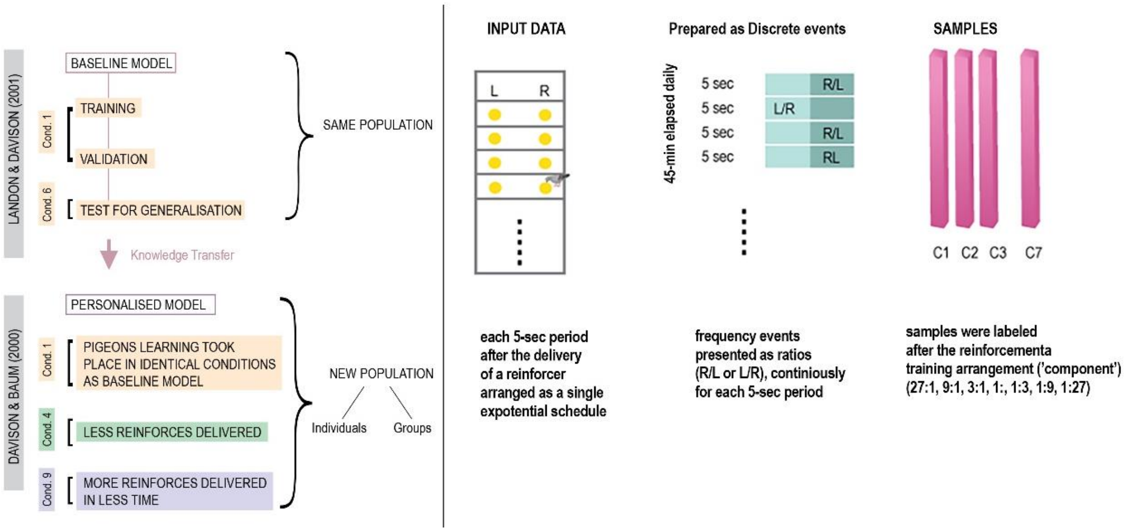

6.2.2 Preparing the data

We prepared the input data following the same procedure as in Plessas et al. (2022) (for more, see 4.2.2). Samples were created by extracting pigeons' choices on the left or the right during a five-second period following the delivery of the reinforcer from all three experimental conditions (Figure 26). We aggregated the data from ten daily sessions in each sample. Each sample consisted of a single component that contained periods of five-second discrete events that occurred in time. The number of periods within each sample varied in each condition, dependent on the number of reinforcers consumed by the pigeons in the prearranged time period - 45-minute daily training sessions. Therefore, each sample consisted of sole pigeons' temporal data and the actual location they occurred in time (left or right). Table 15 presents the number of samples created for each condition to feed the SNN model.

Table 15

Samples Extracted by Pigeons' Temporal Data *(Davison & Baum, 2000)*

Experimental Condition	Samples	Points of 5-sec periods	Total
Cond. 1	24 (7 components*4 samples per pigeon)	36-90 Median: 81 Mode: 90	168
Cond. 4	35 (7 components*5 samples per pigeon)	13-60 Median: 60 Mode: 60	210
Cond. 9	24 (7 components*4 samples per pigeon)	24-60 Median: 60 Mode:60	168

Note: Some samples comprised<90-time points as some of the ten daily sessions at the prearranged time (45 minutes), as the pigeon had not consumed all reinforcers.

6.2.3 Transfer learning to build a personalised single-neuron SNN model.

When creating an accurate, personalised computational model, we use information about an individual and the available information about other individuals related to the same problem. This approach is based on the location of individual dataset samples in relation to the whole space created in the model's architecture (Kasabov, 2019). More on the mathematical computations used in this study in supplementary materials (C5 Appendix C, p.311).

Baseline model

The baseline model was taken from Plessas et al. (2022). The authors set the exact weights, the Leaky-Integrate-Fire (LIF) neuron properties, and the average firing rates of the best-constructed SNN model for each combination of pigeons' data to conduct their generalisation tests. The best-constructed model was the one that used the values of the candidate solution, out of the 40 created, that produced the highest accuracy in classifying the learning history based explicitly on binary choices for all combinations of pigeon data. The best candidate solution, the one with the highest retrodictive outcome, is returned when the training of the SNN with the differential evolution algorithm is completed. These best SNN models are denoted as *baseline* models in this study.

Traditional approach to generalisation

The traditional approach to generalisation requires every new test to be conducted by using the same process as described in the baseline model procedure. This is achieved by testing the performance of the SNN model with a new dataset and calculating the distance of the new firing rates to the mean cluster centroid of the baseline model (vector quantisation method). As with the k-mean algorithm (Han et al., 2012), the vector quantisation method is a simple unsupervised technique for separating the data into groups to divide the number of clusters, by calculating the distance from a cluster's centroid. Therefore, this method aimed to partition i observations (i.e., samples) into k clusters (i.e., seven components), in which each sample belonged to the cluster with the nearest distance to the mean cluster centroid, which is the average firing rate. The distance of each sample to the average firing rates was calculated based on the Euclidean distance (Figure 27).

By using the vector quantisation method, the classification accuracy was measured by calculating the correctly retrodicted learning history, by comparing the new firing rates produced by the new samples fed in the neuron, to the average firing rates created when the baseline model was trained on dataset extracted by Landon and Davison (2001). Then the component label was set (i.e., the label that allows data differentiation for classification) to the reinforcement arrangement with the closest distance (K-nearest neighbour) being k=1 for this study. Finally, the errors were calculated by comparing the actual component the pigeons were trained in with the one that the model retrodicted and obtained the accuracy of the model's performance. The K-nearest neighbour classifier is the least computationally demanding algorithm for supervised learning and one of the ten most effective algorithms in machine learning (Wu et al., 2008). Therefore, the method applied was considered appropriate for the datasets, as it requires no previous knowledge about the data distribution and is also suitable with a small number of variables (Dudani, 1976).

Figure 27

A schematic representation of the KNN (K=1) classifier used in this study when calculating the euclidean distance between the firing rate of the sample and the average firing rates of all components

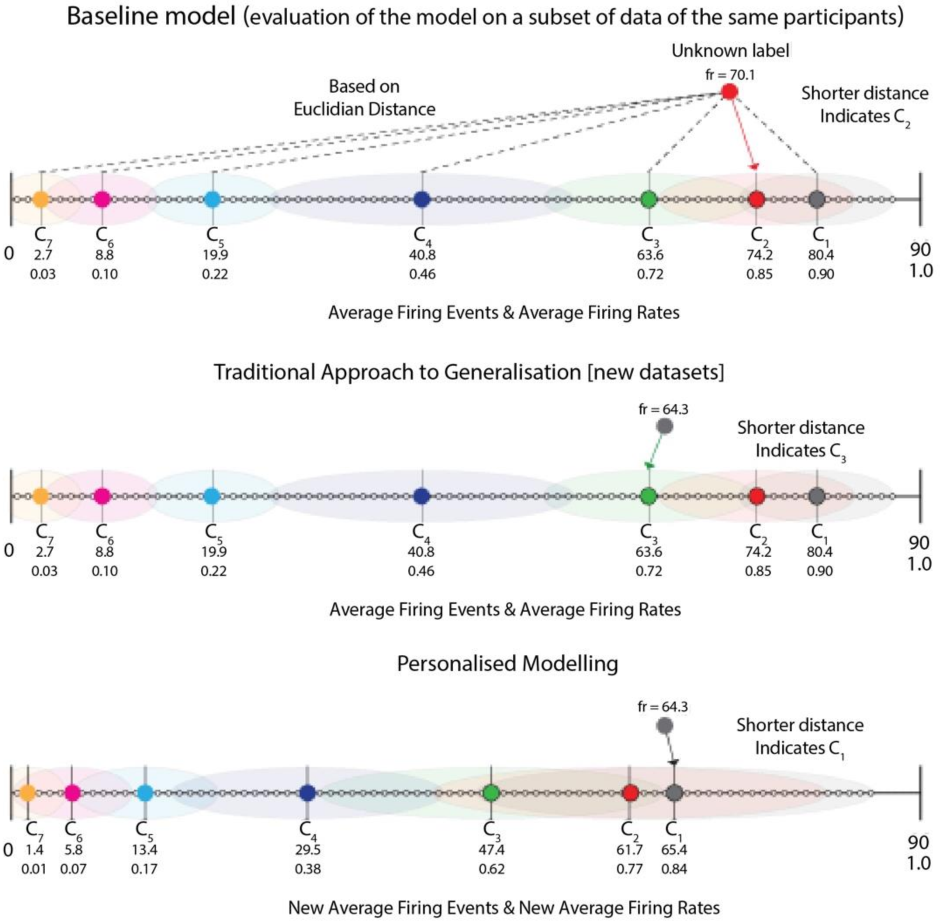

Note. The baseline model refers to the best constructed SNN model, the traditional approach to the generalisation tests based on the baseline model, and the personalised model. The colours on the line represent schematically how the nearest mean cluster centroids changed only for the new model (personalised) as they adapted to the new choice data.

Personalised Model

The approach adopted for this study was based on a transfer learning approach, called network-based deep transfer learning (Tan et al., 2018). As proposed by Tan et al. (2018), the aim was to reuse part of the baseline network including its network structure and connection parameters for the new model, the personalised model. Therefore, in this study, the network parameters of the baseline model— weights threshold and refractory period— were transferred and used with the new datasets to modify the spiking neuron's output by encountering the unique patterns of the new datasets represented through firing rates, produced by the new spike train for each component (i.e., class labels) (Figure 27). Then, for personalised modelling to take place with each new dataset, the average firing rate of the baseline model was not used for classification, being the traditional approach; instead, the new average firing rate produced by the new data (Figure 28). By recalculating the firing rates of the output neurons in an SNN, the frequency of the spikes generated by the neuron was adjusted.

Figure 28

A schematic representation of the artificial SNN architecture for baseline and transfer learning to the personalised approach.

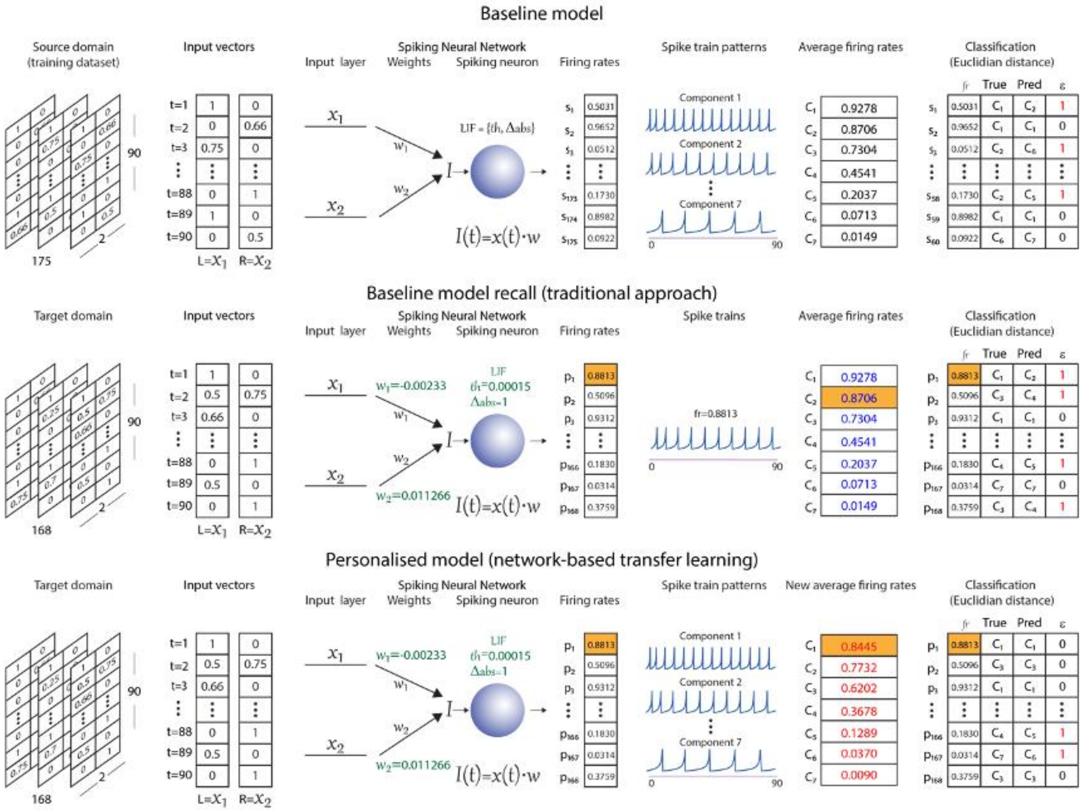

Note. The baseline model refers to the best constructed SNN model. The evaluation of the model is done by having a new set of data that compares the firing rates produces to the average firing rates produced by the baseline model. The traditional approach uses any other new dataset in the same way as the model was evaluated (external validation). However, in a personalised model, the firing rates of new datasets are compared to the average firing rates produced by the new data, rather than the baseline model.

Subsequently, to calculate the new average firing rate that belonged to the same component, the cluster centroids for the k-means algorithm were rearranged so each new sample belonged to the cluster, with the new nearest mean serving as a prototype of the cluster. This process personalised the outcome without further training, as the network remained the same (Figure 27). Therefore, we treated the network as a function and adapted the function of the outputs to identify the learning histories of new participants (pigeons) trained in various environmental conditions.

A summary of the steps followed is in Figure 29.

Figure 29

A summary of the steps for personalised modelling with the SNN

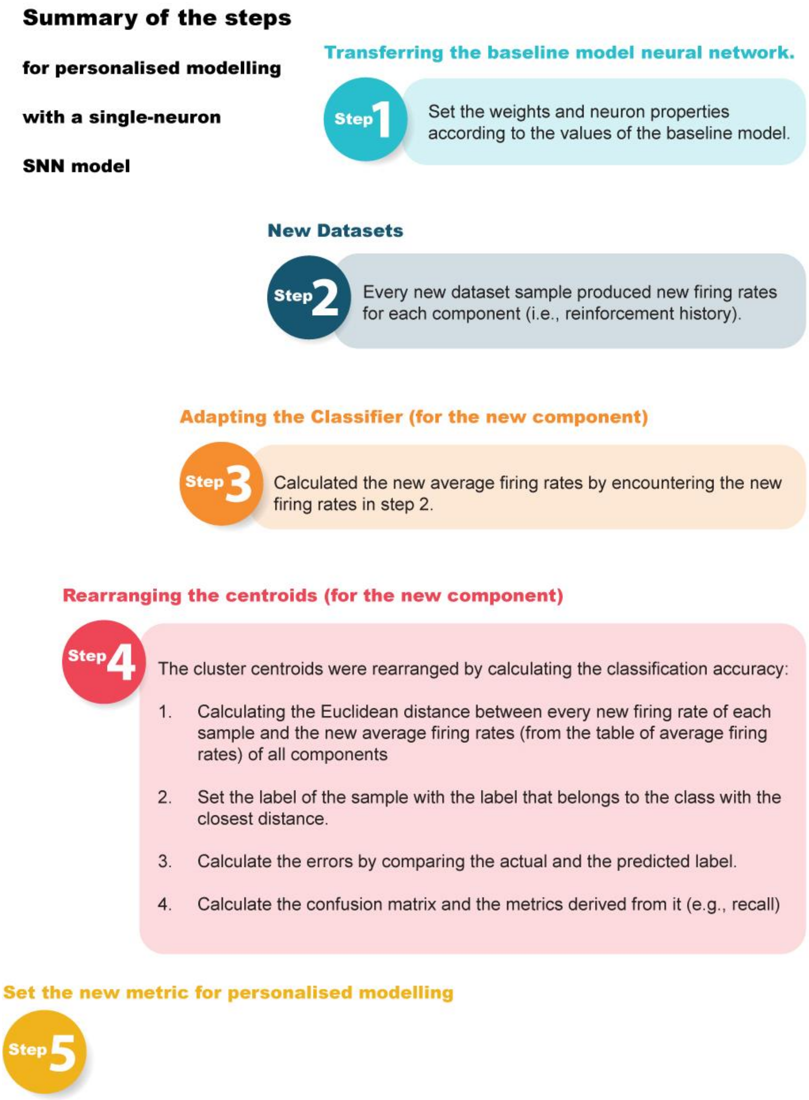

Set the weights and neuron properties according to the values of the baseline model.

Every new dataset sample produced new firing rates for each component (i.e., reinforcement history).

Calculated the new average firing rates by encountering the new firing rates in step 2.

The cluster centroids were rearranged by calculating the classification accuracy:

1. Calculating the Euclidean distance between every new firing rate of each sample and the new average firing rates (from the table of average firing rates) of all components
2. Set the label of the sample with the label that belongs to the class with the closest distance.
3. Calculate the errors by comparing the actual and the predicted label.
4. Calculate the confusion matrix and the metrics derived from it (e.g., recall)

6.2.4 Performance evaluation

We used a series of metrics to evaluate the artificial SNN model's performance with the new dataset. The retrodictive ability was assessed in terms of positive (P) and negative (N) or true (T) and false (F) values. The formulas used for calculations are presented in Table 16 and were based on the output information of the single-neuron SNN personalised models.

Table 16

Performance Metrics used in this Study to Interpret Results and Formulas used for Calculations

Measure	Description	Formula
Accuracy	The fraction of correctly predicted events in relation to all data	$\frac{\# TP}{\# TP + FN + TN + FP}$
Recall (or sensitivity)	The proportion of correctly predicted actual events (i.e.., a true positive) in reference to the total true events	$\frac{\# TP}{\# TP + FN}$
Specificity	The proportion of correctly predicted non-events (i.e.., a true negative) in reference to the total non-events	$\frac{\# TN}{\# TN + FP}$
Precision	The proportion of correctly predicted events in relation to the total number of predicted events	$\frac{\# TP}{\# TP + FP}$
F1-Score	A harmonic mean of precision and recall	
Informedness	A measure about true positives and negatives	$\left[\frac{\# TP}{\# TP + FN} - \frac{\# FP}{FP + TN}\right] - 1$

6.3 Results

Different personalised models were developed to identify the learning history based on the pigeons' binary decision-making. By adopting a network-based approach to transfer the learning from six baseline models developed in Plessas et al. (2022), the firing rates produced by binary choices were redefined to predict these histories, either as a group that participated simultaneously in the same experiment condition or for each individual. The models were tested on six baseline models, denoted as P1-P6 in source study and from now on M1-M6, to allow the evaluation of all best-constructed single-neuron SNN models developed with all combinations of datasets by Plessas et al. (2022). Overall, all personalised models performed better than the traditional approach across all matrices, when detecting a learning history. The results are presented below.

6.3.1 Personalised modelling with different participants datasets across studies

Firstly, the retrodiction of each approach to generalisation (traditional and personalised models) was evaluated with choice behaviours produced in Davison and Baum (2000), experimental condition 1 (Table 15). The dataset extracted from Condition 1 consisted of a small window of other pigeons' binary choices, that occurred contingent upon identical environmental conditions to the datasets used in Plessas et al. (2022), for developing and testing the baseline single-neuron SNN model. Thus, feeding the model with another participant's data allowed us to determine if the traditional approach could generalise and reproduce retrodictions, an external validation assessment, and compare the outcomes with the personalised approach.

Overall, the matrices in Table 17 demonstrate that the ability of each SNN model to generalise with new participants data varied. This outcome was anticipated, as any machine learning model learns from the patterns in the datasets, thus each model's retrodictive ability differs. Nevertheless, the number of all correct retrodictions in relation to the whole dataset, i.e., its accuracy, across all models and with both approaches –traditional and personalised– was good (>70%), with slightly better overall performance than the personalised models. The personalised approach surpassed the traditional when looking into other measures, such as the model's ability to detect the actual learning histories (recall range: 0.63-0.67 against 0.40-0.57) or the trustworthiness of the retrodiction made (precision range: 0.59-0.67 against 0.40-64). In other words, when using the traditional approach with new datasets, approximately half of the time, the models did not detect the false events as false (i.e., that the learning history was other) and the retrodiction was affected by the prevalence of true positives in the dataset. Informedness measures showed that the personalised model was more accurate in making this distinction. These evaluations demonstrate that the spikes generated by the SNN occurred in different times, representing different response patterns in the new dataset in relation to the one used to train the model (Landon and Davison (2001), Condition 1). When these spikes were classified according to the baseline model, average firing rates errors occurred. The data for training the baseline model that produced the average firing rates for the classification task, were not flexible enough to make retrodiction exclusively on the patterns detected in the new dataset.

Table 17

Generalisation Results for Personalised Modelling and Traditional Approach for the Group when using Dataset Extracted by Condition 1, for all Baseline Models Created in Plessas et al. (2022)

Baseline:	Model 1		Model 2		Model 3		Model 4		Model 5		Model 6	
	P	Tr	P	Tr	P	Tr	P	Tr	P	Tr	P	Tr
Accuracy	0.88	0.70	0.86	0.82	0.86	0.83	0.86	0.82	0.86	0.82	0.85	0.80
Recall	0.67	0.40	0.64	0.57	0.64	0.59	0.63	0.57	0.64	0.57	0.63	0.54
Specificity	0.92	0.80	0.92	0.89	0.91	0.90	0.91	0.89	0.91	0.89	0.91	0.87
Precision	0.67	0.40	0.64	0.57	0.64	0.59	0.63	0.57	0.64	0.57	0.63	0.54
F1	0.67	0.40	0.64	0.57	0.64	0.59	0.63	0.57	0.64	0.57	0.63	0.54
Informedness	0.60	0.21	0.56	0.46	0.55	0.49	0.54	0.45	0.55	0.45	0.53	0.41

Note. Personalised models are denoted as 'P' and the traditional approach as 'Tr'

The difference produced by the firing activity of the spiking neuron between the traditional and the personalised approach was further demonstrated in the example represented in Figure 30 when using baseline Model 1. The figure demonstrates schematically that less firing activity was produced for each component when the outputs (i.e., firing rates) were adapted for personalised modelling. Despite the neuron exhibiting less activity, the new firing activity for each component was similar enough to be part of the same component and, at the same time, different enough from other components to discriminate among various components to allow personalisation.

Figure 30

Classification ability for Model 1, as an example of personalised modelling (on the left) versus the baseline model (on the right).

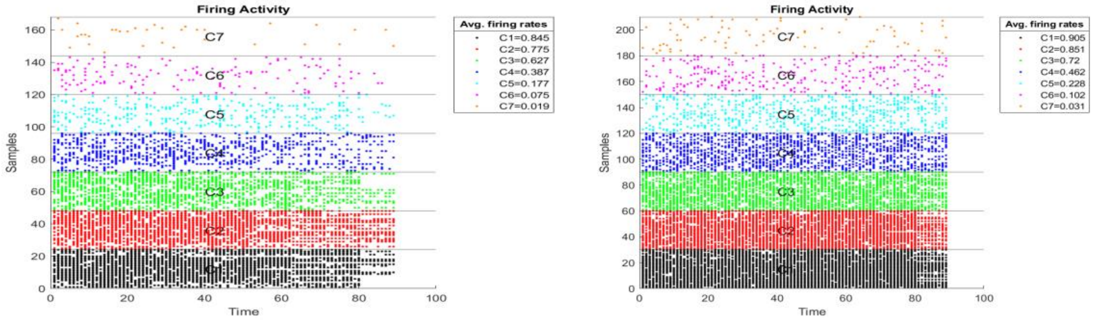

Note: The number of samples differ when calculating the average firing rates table for each model. The dataset extracted from Davison and Baum (2000) study generated 160 samples while the dataset extracted from Landon and Davison (2001) generated 200 samples. This is due to the fact that the former dataset including the last 35 days of training, while the former included 50 sessions of training.

Moreover, in Figure 31, differences are demonstrated in the same example when using baseline Model 1, between personalised modelling and the traditional approach for the dataset extracted from Condition 1 in Davison and Baum (2000). Each figure shows the distribution of firing rates, indicating the probability of a firing rate belonging to a component. For both approaches, there is more firing rate overlap with components closer to each other with the traditional approach to generalisation showed a bigger overlap with choices taking place in environments with similar characteristics (nearby). Consequently, it may be more challenging to identify a distinct pattern for choice behaviour when the choices are similar in proximity to the cluster mean.

Figure 31

An example for Model-1 showing the overlap between components that are closer to each other for with baseline model (on the right) than in personalised model (on the left)

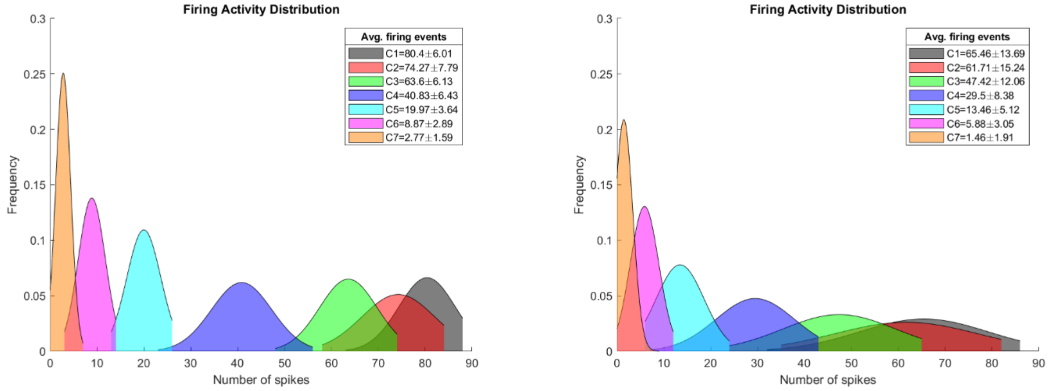

Also, changes in firing rate distribution could vary if the model is fed with datasets for each individual or for a group of individuals, as the retrodiction ability can change. In Table 18, we can view the personalised performance for each participant dataset when tested on each of the six-baseline models in relation to the group dataset. Specifically, for the group and each individual, the accuracy of classifying the component was >=85% for the group and higher for each pigeon (>=90%). In addition, the recall measure was higher for each individual (ranging true positive values between 72%-95%) than when compared to the outcome of a group (62%-65%). Specificity remained high across both approaches (>=92%). This outcome demonstrates that individual differences in choice-making play a role in pattern detection. Therefore, there may be some fundamental learning processes detectable in choice behaviour that may relate to individual differences when each individual responds to changes in the environment. Furthermore, making retrodictions at the group level, can differ when considering the combined diverse patterns from multiple participants.

Table 18

Performance of the Six Models Transferred for Personalised Modelling with data extracted from Condition 1

A S R	M1	M2	M3	M4	M5	M6
P_G	0.88 0.92 0.67	0.86 0.92 0.64	0.86 0.91 0.64	0.86 0.91 0.63	0.86 0.91 0.64	0.86 0.91 0.62
P1	0.91 0.95 0.75	0.93 0.96 0.79	0.93 0.96 0.79	0.93 0.96 0.79	0.90 0.94 0.71	0.90 0.94 0.71
P2	0.95 0.97 0.86	0.94 0.97 0.82	0.94 0.97 0.82	0.94 0.97 0.82	0.91 0.95 0.75	0.91 0.95 0.75
P3	0.96 0.97 0.86	0.96 0.98 0.86	0.96 0.97 0.86	0.94 0.96 0.82	0.96 0.97 0.86	0.95 0.96 0.82
P4	0.93 0.96 0.79	0.91 0.95 0.75	0.91 0.95 0.75	0.91 0.95 0.75	0.93 0.96 0.79	0.93 0.96 0.79
P5	0.95 0.97 0.86	0.94 0.97 0.82	0.94 0.96 0.82	0.94 0.96 0.82	0.97 0.98 0.89	0.97 0.98 0.89
P6	0.92 0.96 0.78	0.91 0.95 0.75	0.86 0.92 0.64	0.91 0.95 0.75	0.88 0.93 0.68	0.88 0.93 0.68

Note. The results show accuracy (A), specificity (S) and recall (R) metrics when trained with all datasets as a group (P_G) and for each individual (P1-P6). The six models used were all pigeons' combinations of the data applied to the baseline model and transferred for personalised modelling.

6.3.2 Personalised modelling with choices with different learning histories

As described in methods, the datasets extracted from experimental Condition 4 and Condition 9 had occurred in environments that differed in specific parameters than those in Condition 1: a) fewer reinforcers were given contingent upon pigeons' choices (n=4) and in Condition 9, the overall speed of environmental change increased (6 reinforcers per minute). As a result, the number, and sizes of samples for each pigeon varied, and the pigeon's choice behaviour changed accordingly to the environmental change (Table 15).

The results in Table 19 demonstrated that personalised models were possible when changes in the learning history applied. In particular, despite the changes, when fewer reinforcers were given (Condition 4 and Condition 9), the accuracy results improved (>88%), in comparison with the performance on Table 16 (Condition 1). The only exception was the performance of the model with Condition 4 with Model-3. Furthermore, the greatest improvement when compared to datasets from Condition 1 is the performance measure of recall. A higher increase in recall (true positives) indicated that the number of reinforcers delivered in each component did not affect the model's ability to discriminate patterns. In other words, varying the sample size from the small dataset had no influence on the retrodictions, and the datasets statistical similarity measures (mean and mode) may play a role (Table 14). Also, when considering the increase in the speed for the environmental change from Condition 4 to Condition 9, recall metrics improvements vary, with slight improvements of Condition 9 over Condition 4, the only exception for datasets used with Model-6.

Table 19

Performance of the Six Models for all Combinations Transferred for Personalised Modelling with data extracted from Condition 4 and Condition 9 in Davison and Baum (2000)

Datasets	Outcomes	M1	M2	M3	M4	M5	M6
Condition 4	Accuracy	0.91	0.90	0.90	0.91	0.91	0.90
	Recall	0.74	0.72	0.71	0.73	0.74	0.73
	Specificity	0.95	0.94	0.94	0.94	0.94	0.94
	Precision	0.74	0.72	0.71	0.73	0.74	0.73
	F1	0.74	0.72	0.71	0.73	0.74	0.73
	Informedness	0.69	0.66	0.65	0.68	0.68	0.67
Condition 9	Accuracy	0.92	0.91	0.92	0.90	0.90	0.88
	Recall	0.76	0.73	0.76	0.73	0.73	0.68
	Specificity	0.95	0.94	0.95	0.94	0.94	0.93
	Precision	0.76	0.73	0.76	0.73	0.73	0.68
	F1	0.76	0.73	0.76	0.73	0.73	0.68
	Informedness	0.70	0.67	0.70	0.67	0.67	0.61

Additionally, as recall metrics improved, when comparing the precision and F1 scores to evaluate the trustworthiness of each model (i.e., precision) created across all datasets (Figure 32), for all experimental conditions, we observed an increase in performances for datasets extracted from Condition 4 and 9 in relation to Condition 1. This implies that the underlying patterns and distributions of the group choice behaviour datasets in Condition 4 and 9 were similar. Thus, the Informedness performance measure improved with both datasets, when compared to performances of datasets extracted by Condition 1, since true positives have been classified better.

Figure 32

Precision scores (same as F-1) on performance for all personalised models with all group datasets extracted by all experimental conditions.

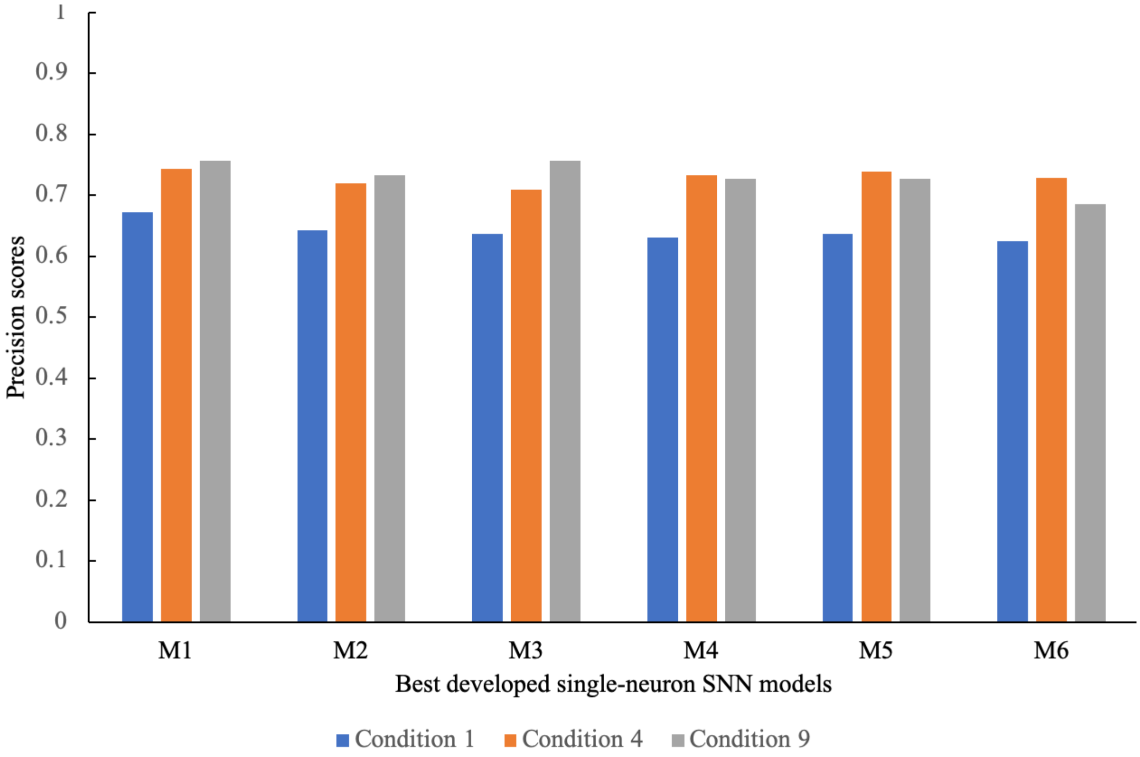

6.4 Discussion

In this study, we demonstrated that using a simple transfer-learning approach, called network-based transfer learning, the SNN model was able to account for individual differences and produce personalised retrodiction based on the new datasets' unique patterns of binary choice behaviour. The retrodictive outcome was achieved by reusing the network parameters pre-trained in Plessas et al. (2022) as a baseline and allowing the artificial neuron to adapt its firing rates to the new datasets.

Consequently, over time the single-neuron SNN integrated the intra- and inter-individual variability of the new datasets to generate its spikes. Then the new firing rates were fundamental for the 'flexible' neuron to produce personalised learning histories. As a result, all datasets extracted from Conditions 1, 4 and 9 in Davison and

Baum (2000) suggested that the model predicted the learning histories with overall accuracy >85%. When the performance of personalised models was compared to the traditional approach to generalisation (Condition 1 dataset), the personalised models performed better. The variability in the data could possibly have affected the spiking activity in the traditional approach. That is, the neuron may have missed features in the data, or the number of inputs altered the spiking activity. However, the latter is unlikely to be the case, as a small number of samples was tested and did not affect the model's ability to generalise (e.g., dataset extracted from Condition 9 and generalisation tests in Plessas et al. (2022). Whereas feature detection in the data (represented as spike-based rate encoding in this study, as reported in Plessas et al. (2022) can be affected by the presence of variability in the data, which in turn influences the production of spikes. The variability in the data proved to be useful for creating personalised models, as it potentially contains valuable information for recalculating the average firing rates. Including this information in the neuron output improved the overall performance of the retrodictive task, and personalisation enhanced the generalisation ability to new datasets.

Furthermore, when recall measures were applied to evaluate the proportion of positive instances in the dataset (i.e., the ability to identify the learning histories that led to the current choice behaviour), the personalised approach improved its performance for the individual and the group. When applied to datasets taken from Condition 1, the model achieved higher recall values, ranging from 64% to 86% in retrodicting the learning history of a single organism, compared to recall values for group performance, which ranged from 62% to 65%. This result makes sense, as personalised models for a single individual can be tailored to the specific response patterns of the individual, whereas group models rely on patterns of the group as a whole. This explains the prevailing view among researchers, that machine learning training requires large data in order to achieve generalisation with any dataset. Yet as Niu et al. (2021) have pointed out, training and testing can have distinctive qualities, hence generalisation is a difficult goal to achieve.

In this study, we demonstrated that large datasets are not necessary for successful generalisation. This research aligns with Ribeiro et al. (2016b) who argues that retrodiction on individual data can inform the development of a "global" model that generalises across a variety of learning histories. For example, personalised models for individuals can help us single out one participant's patterns of choice behaviour and how these affect the outputs for the group in specific directions.

Thus, personalised models for specific individuals can assist us in two ways;

1. to analyse the interpretability of our SNN model by detecting factors affecting the model's firing rates such as environmental change; and
2. to possibly explore learning differences in an individual that is part of a specific group in the relation between learning history and behaviour. Therefore, personalised retrodiction can also hold clinical value, as they can point to learning processes unique to each individual and thus, inform the development of specific environments for learning. So, both individual and group retrodictions are valuable and informative in understanding learning histories.

This modelling approach creates opportunities for future investigations. Further experimentation can use several individual models with a variety of learning histories (e.g., steady state procedures, frequently changing procedures, signalling time.) to expand our understanding of 'universal rules' of adaptive learning and individual level of variation. Given that individual models evolve over time, it would be interesting to investigate how retrodictions are affected by experience as new behavioural data of the same individual are added during successive experiments. The ability of the model to adapt to the data to retrodict and/or the way in which the model fails to do so, may shed light on the specific learning processes required for particular aspects of the environment (e.g., frequency of change, type of consequence).

Therefore, as each performance measure captures different aspects of the model's behaviour, it became apparent in this study, that precision is a metric that fits well with binary choice datasets to capture environmental changes. The trustworthiness of the SNN model in retrodicting an individual's learning history is directly impacted by precision values, which reflect the proportion of positive retrodiction that are actually correct. It may also be used in conjunction with recall and F1 score to evaluate the overall performance of the model when information on both positive and negative examples is important for the detection of learning histories. In this study, precision emerged as a key metric not only for evaluating the accuracy of SNN model retrodiction but also for uncovering the complex relationship between reinforcers and behaviour.

Our results suggest the SNN model was able to detect patterns and retrodict learning histories related to choices made within a 5-second interval from the reinforcer delivery with new datasets. Moreover, the model's retrodictive capability remained consistent regardless of the number of reinforcers delivered in each condition, which was considered in the size and variance of the small sample. This might indicate that pigeon responses may not just be 'preference pulses', that is repetitions of a just-reinforced response; they are the pigeon’s pecking according to where the next reinforcer is more likely to occur. In the analysis presented in Plessas et al. (2022), a small 'window' of observational data that trained the baseline model using log (L/R) response ratios, recording what's emitted in the first 5 seconds following each successive reinforcer delivery, showed the same trend as in Landon and Davison (2001) (see more in Chapter 4). That is, that choice behaviour was shifting according to the reinforcer ratio for each component. When a reinforcer was discriminable, choice-making favoured the key that produced more reinforcement, for example, the 7th reinforcer in Component 7. This showed less discriminable reinforcer ratios within a component (e.g., Component 2); did not produce a solid choice pattern. The SNN retrodiction reflected these patterns. In this study, variations in retrodictions between datasets (individuals and groups), may suggest that pigeons' choice behaviour could be influenced by factors other than the immediate reinforcer delivery, as suggested. If the participants were merely copying the most recently reinforced response, precision scores would have been much higher, reflecting a strict match between pigeon responses and reinforcer deliveries. The results indicate that the SNN model was successful in retrodicting the learning histories of the subjects, even though precision scores were not particularly high. This suggests that the SNN with no previous knowledge on the relation points to factors other than reinforcer delivery, such as the continuation of pecking on the same key producing preference pulses, or the overall reinforcer ratio of the concurrent schedule procedure.

McLean et al. (2014) proposed an alternative interpretation for local analyses of choices postdelivery of a reinforcer. They conducted simulations where they manipulated the means and standard deviations of response distributions in a two-key concurrent variable-interval procedure and observed different preference pulse patterns (i.e., a strong tendency for the first few seconds to the same alternative key). The results in this chapter revealed that variations in behavioural patterns were contingent upon the environmental context in which the reinforcer was acquired. The SNN demonstrated the ability to discern these distinctions, even when they were subtle. If preference pulses are merely an unintended outcome (such as being influenced by a visit interruption caused by the reinforcer), it's improbable for the single-neuron SNN to consistently have achieved precise retrodictions. This was strengthened by the fact the components frequently change, as in the current datasets which made these environments highly unpredictable. Consequently, the information embedded in the datasets encompasses factors that extend beyond mere disruptions of ongoing behaviours. The choice behaviour datasets may be suggesting that the timing (and location) of reinforcer deliveries over time, provided information to the subject about where future reinforcers were likely to occur. Spiking neurons are known to transform data into spikes over a specified temporal window. With time being an inherent feature of spiking activity, it is possible that the firing rates of the single-neuron SNN reflected that the timing of reinforcer deliveries provided information about where future reinforcers were going to be delivered and behaviour tracked this information accordingly (Cowie et al., 2011, 2013). In particular, the SNN model retrodicted on data in the first 5 seconds, often referred as preference pulses, suggesting the model reflected an actual effect of reinforcers on responding when signalled future reinforcer availability (Gomes-Ng et al., 2017). This observation could provide support for the notion that reinforcers are discriminative stimuli signalling future events rather than stimuli solely serving the strengthening of a behaviour even when the contingencies were not explicitly arranged.

A next step of this study is to explore how these findings generalise to other behavioural paradigms. The current study only focused on binary choice behaviour, and it remains to be seen how this model performs on other complex choice tasks (e.g., with three or more choice alternatives). Finally, the study was conducted using data from pigeons, and it is unclear whether the model would perform similarly with data of other species or human subjects. Future studies are needed to address to further validate the usefulness of this approach for modelling reinforcer learning on the paradigm of choice behaviour. Also, specificity measures consistently outperformed recall measures, as in Plessas et al. (2022). This may suggest that the exact weights and the Leaky-Integrate-Fire (LIF) neuron properties that were transferred from the baseline model, reflected patterns that were biased towards identifying true negatives more so than true positive values. In other words, the neuron's experience of detecting non-events was more effective than detecting events. Since there is no solid training method for SNNs, the outcome presented here may be limited to this model's chosen architecture and cannot be generalised to all sorts of other datasets. Further research could investigate the potential influence of the baseline model on the performance of the SNN model, and how this can be optimised for specific research questions. However, from an applied perspective, this does not constitute a barrier for the existing SNN model to be used as a tool for EAB research, as the goal is to have a reliable tool that can make personalised retrodiction on a snapshot of data and not to develop an overarching approach that can solve all problems. What is of interest here, is what personalised modelling can tell us about individual differences when studying learning history and behaviour. In the future, the technique can be tested with other SNN architectures, should they be considered for behavioural research.

In conclusion, the simple and efficient brain-inspired AI tool used here was effective in identifying past learning conditions on the basis of present behaviour. The ability of such a tool to retrodict creates the opportunity to ask new questions in basic research, and hence to understand more deeply the function of reinforcers in controlling behaviour. Further, the tool explored here has potential to shed light on learning history conditions not previously accessible by clinicians. Despite its promise, artificial models are still in its early stages of use in behavioural research. As researchers continue to explore and refine the use of SNNs in studying behaviour through learning histories, we can expect to gain deeper insights into the mechanisms underlying learning and choice behaviour, ultimately leading to more effective interventions and treatments.

6.5 Links between Chapter 6 and Chapter 7

Further, the transferred learning approach utilised in this study, where knowledge applied from one subject can be transferred to another, has provided the opportunity to detect patterns from datasets from different species. This may be helpful to develop learning markers that retrodict future learning patterns; also, to inform on brain processes underpinning learning. It is still unclear if behaviour datasets from various species would still be informative about learning history and which dataset is useful to train the model to produce an accurate retrodictive outcomes. With some preliminary testing (Appendix C) we strongly suspect that approaching the outcomes through multiple species datasets is possible. However, more investigation is required to investigate how this can be approached as research has demonstrated that ecological influences take place in learning (e.g., Breland & Breland, 1961). Lastly, transfer learning approaches are rare with SNN. Future studies should assess the implications of transfer learning with SNN. Improving the robustness of the approach could involve re-training the model using different methods and comparing results across a variety of learning techniques and transfer learning strategies.

Chapter 7 Where do pigeons' choices come from? Transfer Learning in Spiking Neural Networks for Animal Choices in Behavioural Research

Prelude

The ultimate intention around any machine learning model is to train it on a multitude of possible examples, to identify/learn patterns in data. Afterwards, when a new unknown dataset is used as input, the artificial model will be able to identify patterns in the new dataset to make an accurate prediction. The traditional approach to this problem would be to use big datasets to train a machine learning tool to detect any new patterns present in the unknown samples. However, for machine learning to be used with choice behaviour datasets, it needs to be able to handle smaller datasets that represent variability due to individual differences in learning.

With the studies presented in the previous chapters, we showed that an SNN model was able to be trained on small behaviour datasets representing a snapshot of temporal data and to generalise to new datasets. Furthermore, these experiments demonstrated that the outcome was precise and, therefore, trustworthy.

The novelty of this approach lies in demonstrating that a snapshot of observed behaviour carries information about the environment it occurred in, confirming the existence of this relation. Furthermore, it appears that transforming data into spikes is a pivotal factor for generalisation given that the neuron generates new firing rates corresponding to new behaviour datasets.

The goals of this thesis are: a) adding this artificial model to the toolbox of research methodologies, b) translating EAB research by using experimental datasets to answer novel questions, such as 'what is the unknown learning history of the choice behaviour I am observing?', and c) ensuring that observations have both a clinical and basic research value. To reach the overall goal of this thesis, one more step was necessary. We decided to attempt to improve generalisation by using a range of transfer learning methods.

At the commencement of this chapter, an opening section presents a justification of the method of choice to re-train the LIF neuron. This information was not included in the paper presented in this chapter. It should be, and is, included here to show the necessity of refining the method used in the pilot study presented in Chapter 6. The transfer learning methods themselves are described in the paper presented in this Chapter, and therefore, there is no need to expand on them here.

Why nested cross-validation

The goal in this experiment was to run additional experiments with all datasets used in Chapter 6 to assess the single-neuron SNN's robust performance. The term "robust" is used in Cawley and Talbot (2010) to imply "*insensitivity to irrelevant experimental factors, such as the sampling and partitioning of the data to form training, validation and test sets*" (p. 2080).

In the paper presented in Chapter 4, the single-neuron SNN was trained with a *k-fold* cross-validation procedure. First, as described in the paper, the data were randomly split in *k* sets for training (*k-1*) and testing (*1-fold*), where each set had an equal size denoted as *fold*. The split was done repeatedly, and through this process, the classifier algorithm performance was assessed on the tested data. Classification performance was then assessed on a new dataset to get a better measure of the generalisation ability. The method adopted to train the classifier is strongly recommended in the literature. This strategy allowed for an objective, less biased and less optimistic estimation of the model's performance (James et al., 2013). The same single-neuron SNN was used in Chapter 6 to apply a simple transfer-learning technique to allow the neuron to adapt to the new datasets.

In the paper presented in this chapter, we used the same SNN to re-train the model with a computationally expensive procedure known as *nested cross-validation*. Nested cross-validation is a training technique capable of estimating the generalisation error of the underlying model and its hyperparameters search, i.e., the values that determine the learning process of the algorithm and when it will end. Therefore, it reduces the bias in the model but also bias in model selection for best performance in generalisation tests. The extant literature suggests that nested cross-validation can address matters of generalisation error with small datasets showing high variance, to develop a model that can be reproduced and have a generalisable application (Maleki et al., 2020).

The estimation of errors is achieved by using a series of train/validation/test set splits, creating two cross-validation steps. The *k* cross-validation becomes the *outer loop* to assess the final performance of the artificial model for, predicting the reinforcement history. As before, all folds are used as training sets, except for one-fold used once as a test set. Then, each training set is split again randomly into several folds constituting the *inner* cross-validation loop. Thus, the training and testing sets from the inner loop are used to try out different classifier parameters, and the best-performing classifier from the inner loop is then applied to the test set from the outer loop (Kassraian-Fard et al., 2016) (Figure 33).

Figure 33

A Visual representation of a nested cross-validation

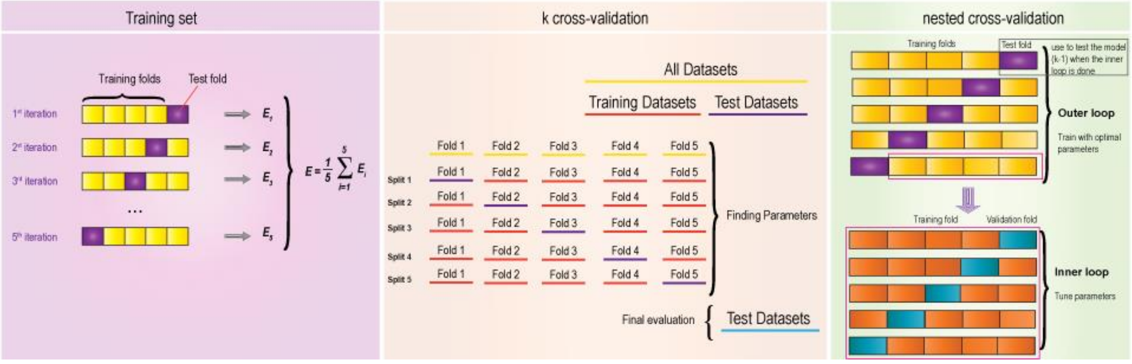

A drawback is that it takes significant time to train the artificial model, as more steps are involved in the process. Hence, it is not always preferred in machine learning training. However, as all datasets used in this Chapter were small, the computational cost was not a particular concern. Therefore, a robust method such as nested cross-validation has the potential to become the method of choice for small behavioural datasets.

Manuscript

Anna Plessas, Josafath I. Espinosa-Ramos, Sarah Cowie, Jason Landon, Dave Parry. "Where do pigeons' choices come from? Transfer Learning in Spiking Neural Networks for Animal Choices in Behavioural Research", 19 March 2023 submitted for review at Neural Computing and Applications.

7.1 Introduction

For the last century, statistical predictive models have been a fundamental standard in psychology (Sarbin, 1944). During this time, researchers have become aware of methodological limitations (excessive reliance on null hypotheses, publication biases, failure to replicate, misinterpretations of p-values), and psychology research has (Imam, 2018; Imam & Frate, 2019; Ioannidis, 2016; Locey, 2020). There is a need, amongst other things, to adopt new methodologies to further psychological science. In this context, machine learning (ML) analysis has been proposed to increase accuracy and improve replicability in psychology research (Orrù et al., 2020).

7.1.1 Prediction in Behavioural Psychology

Prediction and control are predominant goals for behavioural psychology research, a sub-field of psychology. Behaviour psychologists have extensively examined the functional relations between environment and behaviour by manipulating reward arrangements (or reinforcement arrangements) and observing changes in organisms' behaviour. One way these experiments are conducted, is by constantly changing the learning environment for the subject by altering various criteria, such as timing, or the quantum of behaviour needed for the behaviour to produce a reward (Ferster & Skinner, 1957; B. F. Skinner, 1966b; J. Staddon, 2016; Thorndike, 1911). This manipulation allows the researcher to investigate how animals use information from the environment to adjust their behaviour and make choices. This extensive behavioural psychology literature, has inspired some artificial intelligence approaches to reinforcement learning in which algorithms make decisions based on the interaction with the environment, just like a biological agent.

In the last twenty years or so, modelling via ML has seen an increased uptake across scientific domains dealing with complex tasks and practical issues. Trends in psychology suggest that ML tools can enhance our understanding of organisms' behaviour (Yarkoni & Westfall, 2017) and support the objectivity of decisions about clinical treatments (Dehghan et al., 2022; Orrù et al., 2020). However, ML literature in behavioural psychology is scarce (Turgeon & Lanovaz, 2020). One reason is the nature of behavioural-psychology data, usually comprising relatively small datasets of repeated measures of a single organism’s behaviour, that can impair any artificial intelligence model's ability to generalise. Traditional machine learning relies on massive quantities of training data (usually annotated datasets); moreover, training and testing of the model's performance come from the same dataset. In addition, behavioural psychology datasets depicting behavioural changes are diverse and not easily combined because behavioural patterns are shaped by individual differences and the environment. Therefore, training and testing have distinctive qualities and are often dissimilar (Niu et al., 2021) ML model trained on one dataset may not apply well to other datasets. For computer scientists to build machine learning models on behavioural-psychology datasets, it is crucial to consider the characteristics of these datasets and explore how a model can apply to novel instances of the same or similar problem (Barbiero et al., 2020).

This paper addresses the issues of generalisability of behavioural psychology research data by comparing two transfer learning strategies for a spiking neural network (SNN). To the best of our knowledge, this is the first time such an approach has been used with behavioural observations as such. Our results show that transfer learning with a single-neuron spiking neural network (SNN) can serve as a new tool in the psychology research 'toolbox', that could contribute to the methodological limitations the science is facing.

This paper is organised as follows. Section 7.2 starts with a review of previously published work on transfer learning and spiking neural networks and a brief presentation of the behavioural research problem we wanted to solve. Section 7.4 (Methods) starts with the presentation of our preliminary work on the architecture of an SNN model and the training approach we used with choice behaviour datasets consisting of pigeons' right and left choices shaped by different reward arrangments. Further, we present two transfer learning strategies for the SNN we developed. Then we present our case study experiments, applying alternative transfer learning strategies on five published animal learning datasets (Section 7.7). In Section 7.8, the results are compared and discussed. Finally, the section concludes the study with recommendations for future work.

7.2 Related work

7.2.1 Transfer learning in ML

One of the limitations of machine learning strategies is that their capability to make predictions is strictly dependent on the data used for algorithm training. As a result, the algorithm's performance may decrease when the training dataset is too small or when the algorithm is applied to new datasets with some variations (i.e., differences in the quality of the annotated datasets). For example, when a model is trained to classify images of an item during the daytime, its accuracy degrades when it tries to classify a nighttime picture of the same item (Ranaweera & Mahmoud, 2021). Weakness in generalising outcomes implies that some current ML algorithms are ill-suited to recognise familiar patterns in datasets with novel context (e.g., when the domain and feature space of the machine is not the same). An added problem is that retraining a model consumes significant computational resources (for example, by adding more layers, changing learning rates, or adjusting the number of computational units), thereby resulting in extra time and capital investment.

ML transfer learning has shown to be an effective method when the training and testing datasets are drawn from different distributions (Niu et al., 2021). Transfer learning uses a pre-trained model (also called the 'source/baseline model') as a starting point to develop a new artificial model to improve the classification/prediction process (Weiss et al., 2016). Thus, knowledge gained from the baseline model is transferred to the target model to enhance its ability to answer a new problem. A significant advantage of transfer learning is that it reduces the data dependence of the model, thus increasing its ability to generalise to unseen domains. Besides, transfer learning minimises training time, as there is no need to begin from scratch (Tan et al., 2018). Moreover, transfer learning allows us to work better with real-world data that previously were hard for algorithms to manage. Thus, transfer learning unlocks the potential of small datasets and provides for better accuracy by adjusting ML models to new, unseen learning tasks. It suffices to use an existing, pre-trained model as the starting point.

However, transfer learning has yet to be well studied with non-image data in the clinical literature, where a limited number of studies have been published in a recent review(Ebbehoj et al., 2022). There, the authors found that studies primarily used time-series data (n=51, 61%), with fine-tuning being the most common transfer learning strategy (70%), followed by feature representation transfer (22%). Seven studies out of 83 used both approaches. Most predictions represented a binary outcome. To our knowledge, no transfer learning strategies have been reported with discrete events of observable behaviour data, and for ML tools to find applications in behavioural psychology(Dehghan et al., 2022), behaviour datasets are in need of being tested.

Transfer learning has been mainly used with traditional artificial neural networks (ANNs) in deep learning. A variety of ML transfer learning has been proposed to solve real-life applications, including neuroimaging inputs to diagnose diseases (Ardalan & Subbian, 2022); EEG data to verify mental health diagnoses, such as Attention Deficit / Hyperactivity disorder (Lin et al., 2021); patent classification using text labels (Aroyehun et al., 2021); exam scores to improve student performance in higher education (Tsiakmaki et al., 2020); images to assess the quality of fruit and vegetables (Turaev et al., 2020) or wool fibers (Xing et al., 2020); electromyogram-based motion recognition (Hoshino et al., 2022); and vehicle identification (Qian et al., 2022).

7.2.2 Spiking Neural Networks and transfer learning

Although transfer learning has been successfully used in traditional deep-learning ANNs over the last decade or so, research on transfer learning in spiking neural networks (SNNs) is still lacking. The use of transfer learning with deep learning became necessary due to ANNs' tremendous energy consumption and the increase in demand to run models with limited processing capacity and training time (Iman et al., 2023). On the other hand, SNNs, the third generation of artificial neural networks, have gained attention due to their ability to be energy-efficient and the prospect of increasing their efficiency even further in the foreseeable future (Roy et al., 2019).

SNNs are considered one step closer to biological neurons as they transmit information through spikes (or action potentials) by detailed modelling of the neuronal membrane potential and ion conductance (Figure 34). The action potential is a means for the neuron to exchange information; importantly, the neuron’s use of discrete spikes instead of continuous signals, makes it energy-efficient (Tan et al., 2020). Every input to a spiking neuron produces an output spike when the neuron’s threshold is exceeded; if not, the output is zero. Though the SNN architecture is similar to that of traditional ANNs, SNNs can learn from temporal and spatial information. A key component is the timing of spiking neuron models, which mimics how information is transferred in biological neurons (Gerstner & Kistler, 2002). Overall, spikes (or action potentials) from presynaptic neurons will reach a post-synaptic neuron and induce a post-synaptic potential, affecting synaptic characteristics such as travel time or synapse strength. (Ghosh-Dastidar & Adeli, 2009).

Figure 34

A visual representation of an artificial spiking neuron (on the bottom) inspired by a biological neuron (on the top).

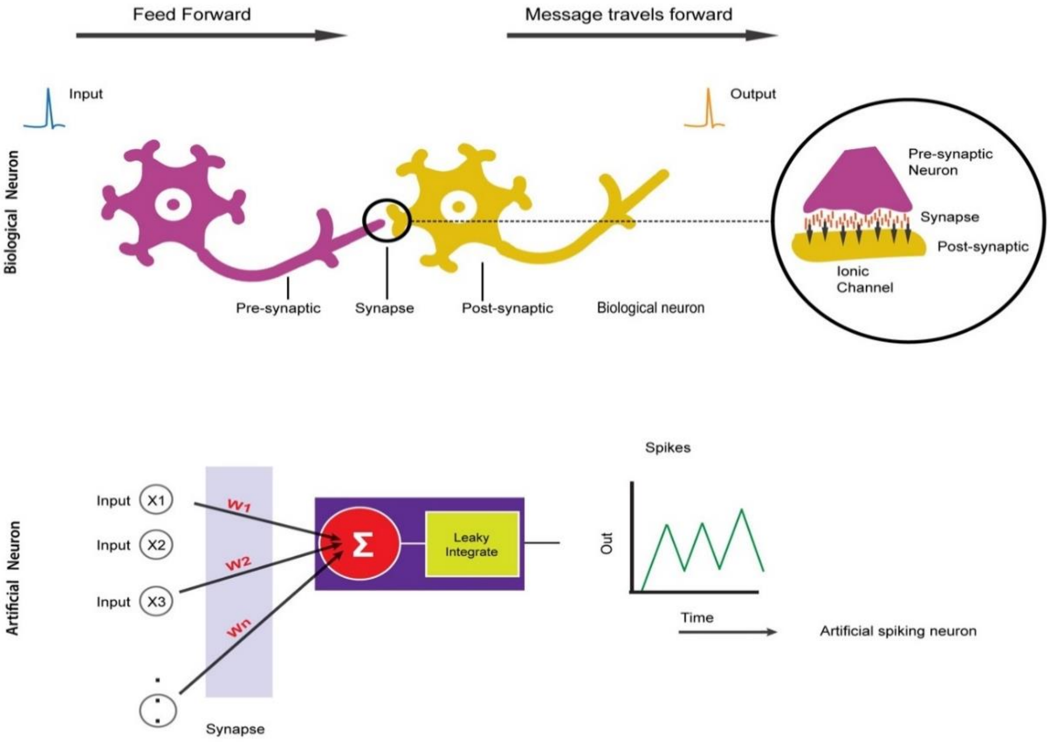

As with traditional ANNs, training SNNs may require large quantities of labeled data that can be costly to obtain (Zhan et al., 2021). Furthermore, as Davidson and Furber (2021) discussed, SNN implementations may prove not to be more energy-efficient than traditional ANNs, as there is a cost involved in accessing the neuron state, and the weights and the choice of encoding can also be detrimental in this process. Thus, they argue that only networks with very low spiking activity justify the use of spikes. Additionally, the common use of many hyperparameters can make the network implementation and optimisation process more expensive. Moreover, the extent of SNNs application can be hindered by difficulties in establishing learning techniques and by SNNs often becoming deeper (Weerasinghe et al., 2021). Therefore, examining transfer learning approaches for SNNs is important, as they promise to overcome issues of costs in computational energy and resources and have the advantage of being more biologically plausible.

A recent study examined using SNNs with transfer learning for object recognition by using layer mitigation (Fu & Dong, 2021). Based on the initially trained network, some layers were fixed without retraining the model for the new task. As a result, the accuracy of classifying their datasets was high (>=93%), and the SNN achieved better results than the benchmarking approaches, suggesting that transfer learning was fundamental for the SNN to improve its performance on different tasks (Fu & Dong, 2021). Transfer learning has also been used with deep SNNs, with the insertion of a domain loss after the feature extraction layer and by optimising the network at the same time for the domain loss and the classification task (Zhan et al., 2021). Another popular approach has been to train a CNN network first by using backpropagation and then by transferring the weights to the target SNN model to minimise the loss (Yan et al., 2021). Overall, deep learning with spiking neurons has shown promising results and increased real-world applications and is becoming competitive to conventional ANNs (Pfeiffer & Pfeil, 2018), especially as they are overcoming some of the abovementioned issues.

This paper discusses the application of transfer learning on a single-neuron spiking neural network architecture with data from behavioural observations (non-image or physiological data). In (Beniaguev et al., 2021) they investigated the computing power of an individual neuron. Their research was premised on the neuroscientific idea that single neurons have several functional subregions, as their complex branching implies that single neurons might behave as extensive networks. Furthermore, they suggested that networks formed by single neurons have a unique architecture supporting their computational power. In an earlier paper (Plessas et al., 2022), data from behavioural research experiments with pigeons were extracted to train a single-neuron SNN, and the classification of pigeons' past learning histories (unknown to the machine) was quite accurate when pigeons' choices served as the only input to the spiking neuron. This study examines whether transfer learning strategies can be used with a single-neuron SNN architecture processing behavioural and temporal variables such as pigeon choices.

7.2.3 Experimental analysis of behaviour and behaviour change

Adaptive learning is a fundamental process that defines living organisms. Psychology and neuroscience study human and other living organisms' choices using a wide range of behavioural, neuroimaging and cognitive experiments to understand how learning develops. Learning experiments following experimental analysis of behaviour (EAB) procedures, to explore what environmental factors affect choice over time (see Cowie & Davison, 2016a for a review). A reward obtained in a particular context over successive trials can change behaviour (learning) (Skinner, 1938; Thorndike, 1911). In choice experiments, organisms choose between two or more options, each of which is linked to a reward/ punisher based on a defined rule set for each option. The defined rules (i.e., reward arrangement) determine how often the reward/punisher will be delivered based on how the individual allocates their responses or based on the time when other options are available. The experimenter observes how the organism uses the information in the environment (i.e., the reward arrangement) to make choices and how these choices change over time. These studies generally occur in laboratories with animals. Mathematical models have shown that animal choice behaviour is orderly (Baum, 1974; Davison & Jenkins, 1985) and conforms to the same general rules as human choice (Cowie et al., 2021; Podlesnik et al., 2022). The relationship between choice and environment can be described by linear regression by analyzing choice data as logarithmic ratios of the subject responses and ratios of the reward obtained (Davison & McCarthy, 1988).

While mathematical models applied in behavioural psychology use the obtained rewards in an attempt to predict choice, they tend not to permit the reverse - predicting obtained rewards on the basis of choice. ML provides a unique advantage here. Recent work by the current authors has used ML techniques to explore predictive relationships between choices and reward arrangments (Plessas et al., 2022). A single neuron SNN was trained and achieved high performance across metrics (accuracy ≥93%, specificity ≥96%), precision ≥81%) and F1≥81%). The authors showed that using choice behaviour as a temporal variable could provide information about the training pigeons received in the past (i.e., the reward arrangments) when this past is unknown and what only is currently observed, is available. The SNN model prediction was based on a small snapshot of data (i.e., not the entire dataset), using pigeons' left and right choices as input vectors. It was also able to generalise to another dataset for the same population. However, variability in organisms' behaviour means their choices are adapted to the environment, creating unique patterns that can differ due to individual differences. Therefore, different datasets or different amounts of data may alter the model's ability to generalise. The number of choice responses and their timing could be critical in setting a learning 'marker' to make predictions about reward arrangments. Machine learning models can improve their predictions based on experience, thus contributing to expanding researchers' behavioural knowledge, from existing experimental studies to predictions of organisms' learning ability. Furthermore, predictions in choice behavioural research based on ML can contribute to our decision-making process in determining appropriate therapeutic interventions. Therefore, ML algorithms can find a utility in behavioural and clinical sciences, by providing means to infer future behaviours. As encouraged by recent research discussions (Dehghan et al., 2022) this may be accomplished using behavioural data.

7.3 Main contributions of this study

In this study, we developed a single spiking neuron (SNN) architecture for creating baseline models of two pigeon populations trained to make binary choices in similar environments (i.e., identical and other reward arrangments). We compare the models' performance (accuracy, precision, recall, and F1) and transfer learning from one model to the other and vice versa. We propose two transfer learning strategies: feature extraction and fine-tuning. Both strategies have the potential to be applied when SNN is implemented for behavioural research allowing the SNN model to deal with small variable behavioural datasets. Transfer learning makes the model more flexible and can, therefore, assist behavioural scientists in improving research on decision-making processes when designing reward arrangments to change choice behaviour.

7.4 Method

This section presents the architecture for the developed single-neuron SNN. First, we will show how SNN model training and generalisation tests were conducted using pigeons’ binary choices to predict the reward arrangement in effect. Then we will present feature extraction and fine-tuning as transfer learning strategies for developing a new model for new datasets.

7.5 Single spiking neuron network for the baseline model

7.5.1 SNN architecture

We built an SNN based on the architecture described in (Vazquez & Cachón, 2010) and extended it to recognise temporal data (i.e., pigeon's binary choices). The architecture consisted of five main elements:

1. A computational unit based on the Leaky Integrate-and-Fire (LIF) neuron model (Gerstner & Kistler, 2002).

   $$\tau_m \frac{dv}{dt} = -v(t) + RI(t) \quad (1)$$

   *Where $v(t)$ represents the membrane potential at time t, τ_m is the membrane time constant, R is the resistance, and $I(t)$ is the input current at time t. When the membrane potential $v(t)$ reaches a threshold value th, it emits a spike and resets to a reset potential v_r. The neuron stays inactive during an absolute refractory period Δ_{abs}. Then, the leaky integration process is re-initiated following a delay of Δ_{abs} after the spike.
2. A set of input patterns $p = \{X^i, k\}_{i=1}$ where X^i is an $n\text{-}by\text{-}2$ matrix and each column vector $x_1^i, x_2^i \in \mathbb{R}^n$ contains the n pigeon's choices (either to the left x_1^i or to the right x_2^i key), and $k = 1,2,\dots,K$ is the class (class=reward arrangement or briefly component) to which X^i belongs.
3. A vector of weights $w \in \mathbb{R}^2$ that represents the "synaptic connections" between the input nodes and the neuron.
4. An input signal $I_i = X^i \cdot w$ that simulates an injected current that feeds the neuron to produce a spike train $s \in \{0,1\}^n$ where 1 represents a spike (action potential) and 0 is no spike.
5. Lastly, for classification, we first calculated the firing rate of each input pattern as,

   $$fr_i = \frac{spk}{|s_i|} \quad (2)$$

where spk is the number of spikes. Then we calculated the average firing rate $F \in \mathbb{R}^k$ of all the samples belonging to each class k. Then, we determined the class of an input pattern X^i as the label of the class with the closest average firing rate using Equation (3)

$$c_i = arg\, min_{k=1}^{K}(|F_k - fr_i|) \quad (3)$$

7.5.2 SNN training

We addressed the training as an optimisation process, where the objective function was to find the vector of weights w that maximises the classification accuracy of the input patterns p. In this study, we applied the Differential Evolution (DE) algorithm (DE/rand/1/bin), a population-based heuristic method for global optimisation over continuous spaces (Storn & Price, 1997).

To exert maximum leverage of the algorithm, we also included the LIF neuron's threshold th and absolute refractory time Δ_{abs} as elements of the candidate solution, the DE population q. These two properties of the neuron affect the output spike train and the firing rate used for the classification. In this study, we define a population of candidate solutions (agents) as $q = \{a_1, a_2, \dots, a_{NP}\}\ \forall\ a_i = \{w_1, w_2, th, \Delta_{abs}\}$ where $NP = 10d$ is the population size and d the size of the agents (dimensionality of the problem). A representation of the single neuron architecture and the DE agent is shown in Figure 35. The models created using the proposed architecture and training strategy are referred to as baseline models.

Figure 35

A schematic representation of the product of an input and weights that produce an injected current that feeds the spiking neuron to produce a spike train with a specific firing rate.

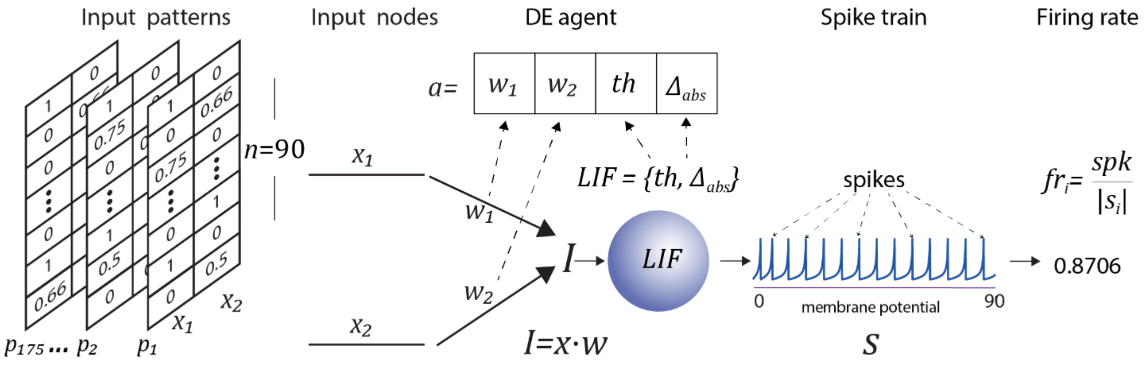

7.6 Transfer learning for SNN

We performed transfer learning similarly to how we applied DE to train and create baseline models. Initially, we allowed the $u_{base} = \{w \in \mathbb{R}^2, th, \Delta_{abs}\}$ to be the underlying patterns of a pre-trained (baseline) model, p' the dataset of input patterns of a new source domain and $F' \in \mathbb{R}^k$ the outputs (average firing rates) of a new SNN model. First, we transferred u_{base} to a new SNN. Then, we trained it on a testing dataset $p'_{train} \subset p'$. Finally, we generated new firing rates F' that we used to classify a testing dataset $p'_{test} \subset p'$.

In our study, we propose two approaches to transfer learning: feature extraction and fine-tuning.

7.6.1 Feature extraction

In this modality, u_{base} was transferred to a new SNN. Then, the outputs F' of the SNN were adjusted by re-calculating the firing rates when using a training set p'_{train}. The performance of the model was assessed using Equation (3) to classify the testing dataset p'_{test}. Note that F_k in Equation three must be replaced by F'_k. Figure 36 shows a comparative diagram of a baseline model and a model created using transfer learning.

Figure 36

The weights and LIF parameters (in green) of the baseline model were used to train a new SSN on a new dataset.

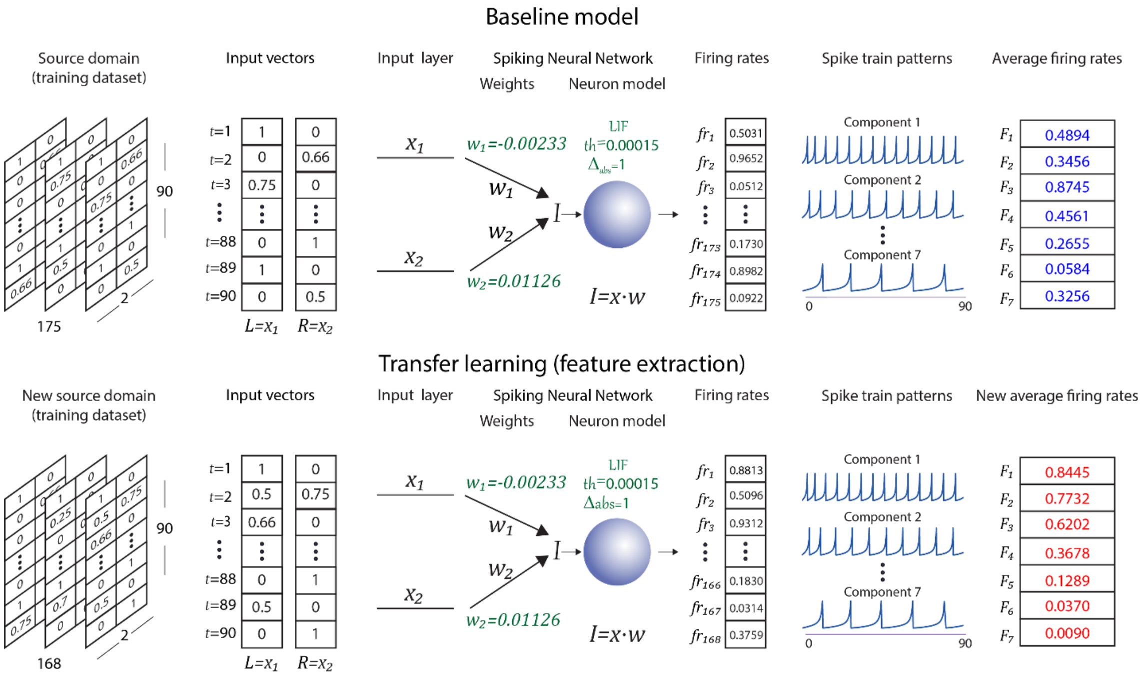

Input patterns from different domains produced different average firing rates (in red), which were more suitable to the problem.

7.6.2 Fine-tuning

In fine-tuning, a set of underlying patterns U is transferred to a new SNN and evolved through DE (fine-tuning) to find the new firing rates F' that best classify a training dataset p'_{train}.

Let $q'_{init} = \{U \ \cup A\}$ be the initial DE population used for training the new SNN, $U = \{a_1, a_{2,} \dots\}$ the set of underlying patterns where a_i is the fittest agent of the last DE population of a baseline model q_{base_i}, $i: i \leq NP = \{1,2, \dots\}$, or $U = \{a'_1, a'_2, \dots\}$ where $a'_i \in q'_{base_i} : q'_{base_i} \subseteq q_{base_i}$, and finally, $A = \{a'_1, a'_2, \dots, a'_m\}: a'_i \in \mathbb{R}^4$ is a set of new agents randomly initialised, and $m: m = NP - |U'|$. The set A can be formed using the following approaches:

1. Uniform: A random set of agents $a \in \mathbb{R}^4$ with a uniform distribution.
2. Bounded: A set of agents where each of its elements $\{w_i, w_2, th, \Delta_{abs}\}$ is a random number with uniform distribution subjected to constraints previously defined by the user e.g., $th = [0.01, 0.05]$.
3. Normal: The agents were created with a normal distribution based on the mean (μ) and standard deviation (σ) of the underlying patterns U.

We hypothesised that a predefined search space could help the DE find a better solution in fewer iterations than an initial search space created with an entire population of random agents. In Figure 37, we show a representation of the fine-tuning approach.

Figure 37

The underlying patterns (green) of a baseline model serve as part of the initial population of the Differential Evolution of the new SNN.

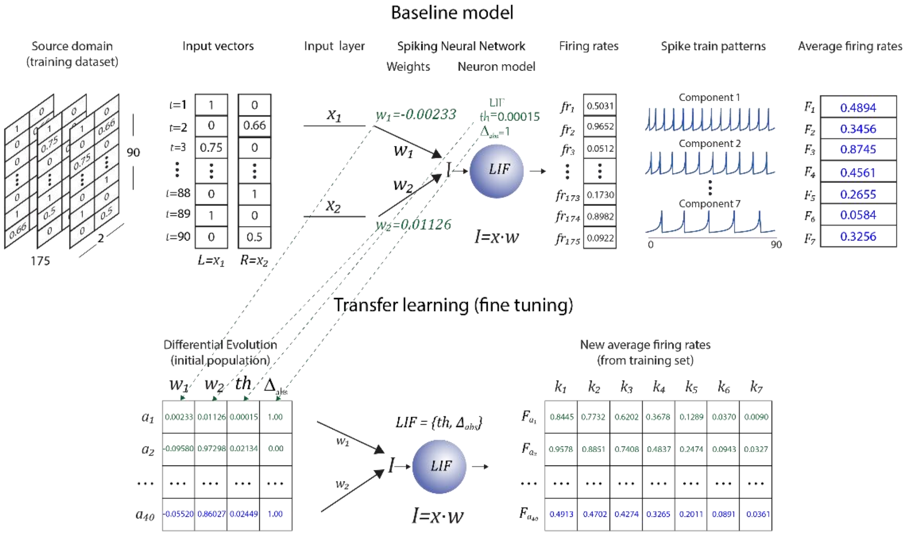

Note. The remaining agents (set A in blue) are randomly initialised following one of the three approaches. Every agent of the population produces different average firing rates Fa_i for each class k.

7.7 Experiments/case studies

7.7.1 Datasets

A reward arrangement is any procedure that manipulates rewards (or *reinforcers*) based on a defined rule; the reward is food for pigeons, and the *arrangement* is the rule used to deliver the reinforcer (J. E. R. Staddon & Cerutti, 2003). We used data from several experiments with pigeons as participants, published in two seminal behavioural psychology papers (Davison & Baum, 2000; Landon & Davison, 2001) that demonstrated how choice behaviour changes when the reward arrangement changes (i.e., component). Their sophisticated procedures showed that participants shifted their choices in relation to the arranged reinforcer regime and, depending on the arrangement, the rate of behaviour change varied; therefore, learning patterns differed for each component and each pigeon.

The pigeons were housed individually in their home cages, which served as their experimental chamber (Figure 38).

Figure 38

Pigeons' experimental chamber in behavioural research laboratories

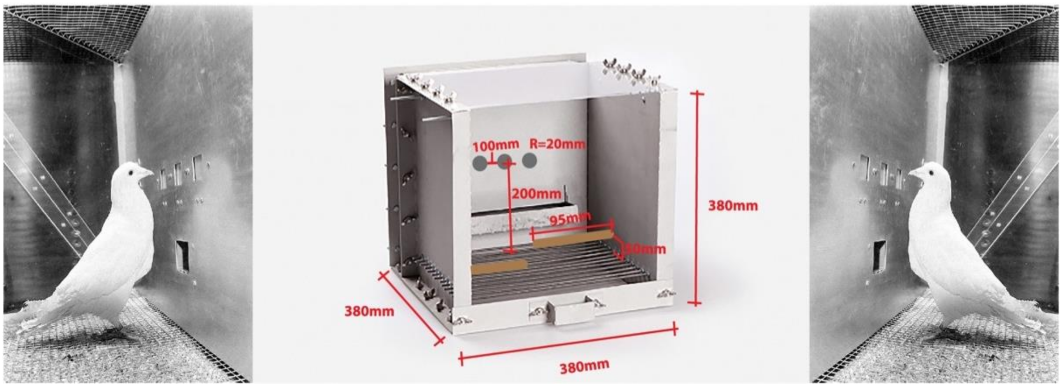

Pigeon responses to the illuminated keys exceeding 0.1N were recorded. A reward arrangement was set for each key and changed rapidly and randomly according to a prearranged rule set by the experimenters (Figure 39). Each training session was organised by setting seven different components with (L[left]: Right[right]) reinforcer ratios as 27:1, 9:1, 3:1, 1:1, 1:3, 1:9, 1:27. Across experiments, these reward arrangments were either identical or differed in rule settings parameters, though presented randomly for each session (Table 20).

Figure 39

A visual sample of a pigeon making choices during experiments.

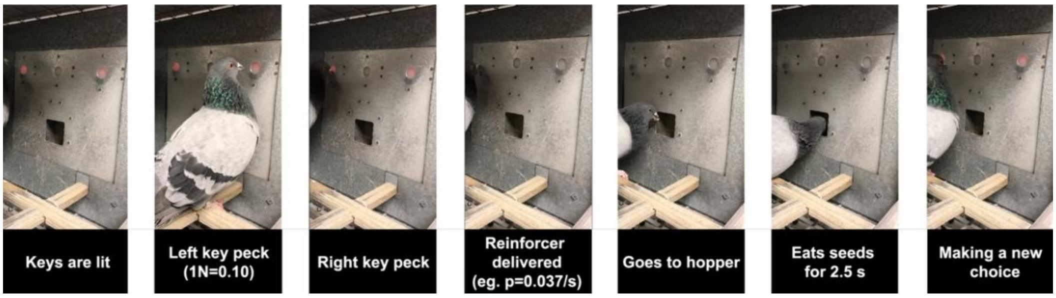

Table 20

The *Characteristics* of The Datasets used for SNN Prediction

Study	Pigeons	Experiments	Reward arrangments	Daily Sessions	Reinforcer per minute	No of reinforcers
L & D (2001)	P61-P66	Cond. 1 (A)	7	50	2.22	10
		Cond. 6 (B)	7	35	2.22	10
D & B (2000)	P91-P96	Cond. 1 (C)	7	35	2.22	10
		Cond. 4 (D)	7	35	2.22	4
		Cond. 9 (E)	7	35	1.00	6

Note: Pigeons 61-66 were exposed to Condition 6 a long time after their exposure to Condition 1; in the interim, they were trained under Conditions 2 to 5, each of which was in effect for approximately 50 days. The same applies to Pigeons 91-96.

7.7.2 Preparing the data

We prepared our temporal data applying the same extraction rule as in (Plessas et al., 2022), to allow the results to become comparable when dealing with a snapshot of data rather than a complete dataset of response patterns. Therefore, we only extracted pigeon responses as a discrete frequency event that occurred continuously for five seconds, immediately after the delivery of the reinforcer. The input vectors created to train the SNN model were the final outcome of these behaviour research experiments- the choice response of the pigeon by pecking to the left (L) or to the right (R) key. The output of our SNN model was to predict in which environment (i.e., reward arrangement) those choices were shaped when the reward arrangement was unknown (Figure 40).

Figure 40

Schematic representation of how we approached choice data for training the SNN model.

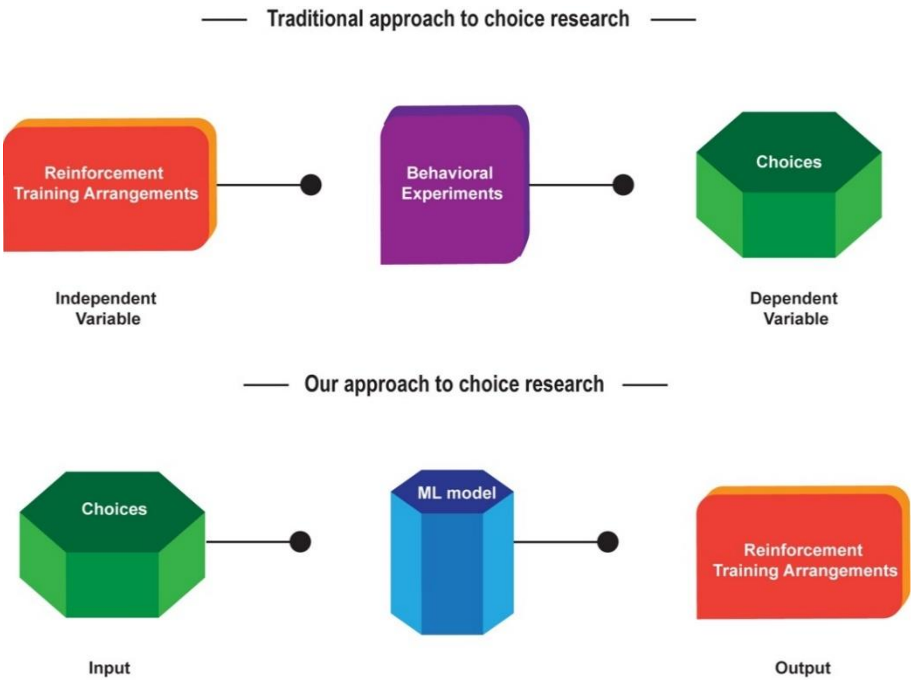

Note. The top row indicates the traditional choice research approach, and our approach is at the bottom.

Samples were created by extracting pigeons' choices on the left or the right during a five-second period following the delivery of the reinforcer from all five experimental conditions from the two published papers (Figure 41). We aggregated the data for each sample from ten pigeon training sessions. Each sample consisted of a single reward arrangement (component) that contained periods of five-second discrete events that occurred in time. The number of periods within each sample varied in each experimental condition that was dependent on the number of rewards consumed by the pigeons in the prearranged time (45-minute daily training sessions). Table 21 presents the number of samples created for each condition to feed the SNN model.

Table 21

Samples Extracted by *Pigeons'* Temporal Data

Experimental Condition/ Datasets	Samples	Points of 5-sec periods	Total
(A)	35 (7 components*5 samples per pigeon)	81-90 (9 reinforcers*10 days)	210
(B)	21 (7 components*3 samples per pigeon)	81-90 (9 reinforcers*10 days)	126
(C)	24 (7 components*4 samples per pigeon)	31-90 (9 reinforcers*10 days)	144
(D)	35 (7 components*5 samples per pigeon)	51-60 (6 reinforcers*10 days)	210
(E)	24 (7 components*4 samples per pigeon)	51-60 (6 reinforcers*10 days)	144

Note: Some samples consisted of >90 points as some of the ten daily sessions ended at the prearranged time (45 minutes), and the pigeon had not consumed all reinforcers.

Figure 41

Datasets extracted from two source papers for SNN modelling.

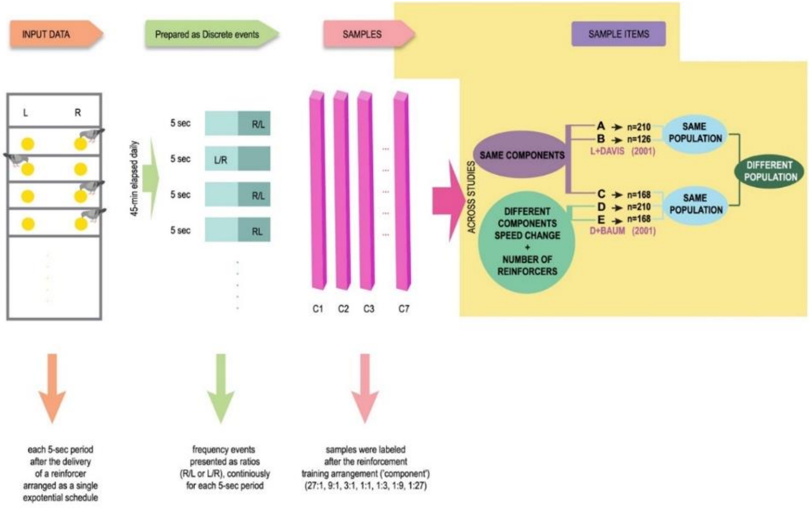

7.7.3 Procedure for training SNN models

All five datasets (Table 21) were used to create baseline models. The baseline SNN models created from one population were used to apply the transfer learning with the datasets that did not participate in the training and represented either a different population or a different reinforcement training arrangement. Ten baseline SNN models were created for each dataset. As described in the method section (7.5.2), nested cross-validation (10 outer and 5 inner stratified k-fold cross-validation) was applied for hyperparameter optimisation and model selection. After training the neuron with Dataset A (i.e., Model_A), we classified Dataset B using the baseline models. In training the SNN model, the weights, neuron parameters, and firing rates of the baseline model were used for classification. Therefore, no calculation of new firing rates was needed. Taking a step further, we classified Dataset B by applying the feature extraction modality of transfer learning. As described in the method section (7.5.1-7.5.2), the weights and neuron parameters remained fixed for these experiments, but the new firing rates were calculated from the dataset. Datasets C, D, and E were also used with the baseline models and then applied for both transfer learning strategies. We then conducted the reverse experiments by switching the datasets, allowing all models' performances to be compared (Figure 42).

We used a series of metrics to evaluate the performance of the artificial SNN model for all datasets. The results were assessed in terms of positive (P) and negative (N) or true (T) and false (F) values to calculate precision, recall, accuracy, and F1-scores.

Figure 42

Schematic representation of baseline and transfer model training.

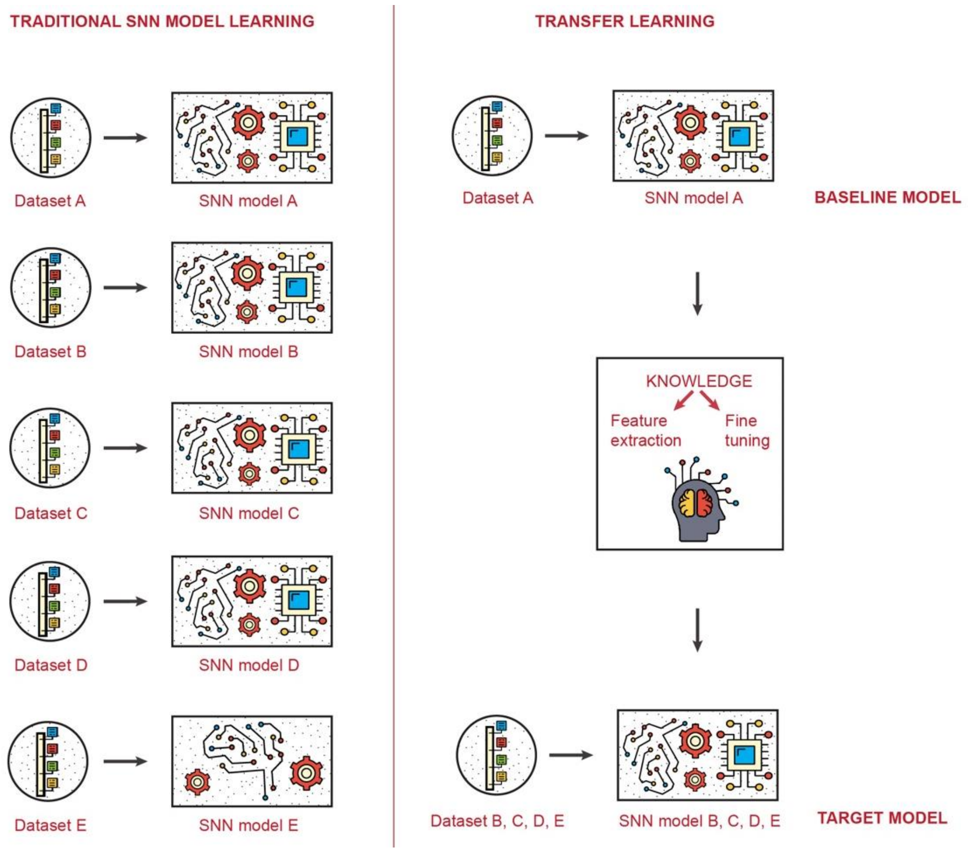

Baseline Model

We used the parameters described in Table 22 for training the single SNN. We implemented a nested cross-validation procedure to estimate the performance of the model. We also used the same procedure first for optimising the hyperparameters (underlying parameters u_{base}) of the model on a dataset (p') and then for comparing and selecting a model of the dataset. We used the stratified k-fold cross-validation strategy in both the outer and inner loops and $k_{outer} = 10$ and $k_{inner} = 5$, respectively.

Table 22

Differential Evolution (DE) Parameters

Description	Variable	Value
Strategy		de/rand/2/bin
Dimensionality	d	4
Population size	$NP = 10d$	40
Crossover rate (recombination)	Cr	0.7
Differential weight (mutation)	f	0.1
Weight bounds	w	[-0.01, 0.01]
LIF threshold bounds	th	[0.01-0.5]
LIF Absolute refractory period bounds	Δ_{abs}	[1-10]
Maximum number of iterations	$maxiter$	300

Feature extraction

We only applied the stratified k-fold cross-validation strategy for the experiments with $k = 10$. No hyperparameter optimisation was needed, but the calculation of new firing rates was applied by using Equation (3) on every dataset for training.

Fine-tuning

In our experiments, the initial DE population to train the new SNN was $q'_{init} = \{U \cup A\}$ being $U = \{q^1_{best}, q^2_{best}, \ldots, q^{10}_{best}\}$, where q^i_{best} is the fittest agent of the $i-th$ baseline model that the outer loop of the nested cross-fold validation produced, and A a set of 30 agents randomly initialised according to the approaches described in 7.6.2Section 3.2.2. Table 23 shows the DE parameters used in fine-tuning transfer learning. Further, to estimate the model's performance, we implemented a nested cross-validation procedure with the same configuration described in section 7.5.1.

Table 23

Differential Evolution Parameters for Fine-tuning. Parameters w, th and Δ_{abs} are used in the Bounded Approach for Creating the DE Population.

Description	Variable	Value
Strategy		de/rand/2/bin
Dimensionality	d	4
Population size	$NP = 10d$	40
Crossover rate (recombination)	Cr	0.7
Differential weight (mutation)	f	0.1
Weight bounds	w	[-0.01, 0.01]
LIF threshold bounds	th	[0.01-0.5]
LIF Absolute refractory period bounds	Δ_{abs}	[1-10]
Maximum number of iterations	$maxiter$	300

7.8 Results and discussion

To present the performance of the models, we combined the F1-score of each class to create a single measure for the whole model (micro-averaged F1-score). By this, we provided a global measure of our performance and all the measures (i.e., accuracy, recall, precision, micro-F1) became equal. We also calculated the weighted-average F1-score to present the weighted mean of the measures. The weights are the total number of samples of that class – the number of samples created for each component, for each pigeon. As we used balanced datasets for our multiclassification problem, negligible variations in our F1-score were denoted after the second decimal. Therefore, we presented only the weighted F1-score for the case studies as the results were equal.

7.8.1 Training baseline SNN model and validating each model's performance

The performance of all SNN models was evaluated (Table 24). Five different models were developed when using five different datasets as inputs. All models were capable of handling the different datasets, with Model$_A$ and Model$_B$ (i.e., SNN models trained with Datasets A and B) performed the best with a mean score of 83%.

Table 24

Results of Weighted F1-score FOR BASELINE SNN Models Without Transfer Learning

Models	1	2	3	4	5	6	7	8	9	10	Mean	SD
Model$_A$	0.90	0.81	0.86	0.86	0.76	0.71	0.81	0.81	0.86	0.90	0.83	0.06
Model$_B$	0.92	0.85	0.92	0.92	0.92	0.77	0.67	0.75	0.83	0.75	0.83	0.09
Model$_C$	0.65	0.65	0.65	0.76	0.71	0.65	0.65	0.65	0.69	0.75	0.68	0.04
Model$_D$	0.67	0.57	0.81	0.62	0.71	0.76	0.81	0.67	0.76	0.76	0.71	0.08
Model$_E$	0.82	0.76	0.82	0.82	0.76	0.94	0.76	0.59	0.75	0.81	0.79	0.08

Note: Number of iterations (n=300). The name of each SNN model created corresponds to the dataset used for training the model with nested cross-validation.

To validate the performance of each developed SNN model, we ran several generalisation tests with datasets that did not participate in the training, to assess each model's prediction ability when no transfer learning was applied. We used two case examples for generalisation.

In the first scenario, we assessed generalisability with SNN Model$_A$ and Model$_C$. The models' were created with datasets from different populations, yet the input data came from identical behavioural experiments (i.e., all participants' choices were shaped under identical reinforcement training arrangments). Generalisation tests were then run with the remaining four datasets (Table 20). Overall, generalisation tests were successful, and performance accuracy ranged between 64%-84% in identifying the component. In particular, the prediction performance was superior to all other generalisation tests with Dataset B with Model$_A$ and Dataset A with Model$_C$ (Table 25). Therefore, when using the pigeon's choice to predict the training environment, a traditional machine learning approach (i.e., to first train each model and then test for generalisation) can be effective even with different populations participating in identical behavioural experiments.

In the second scenario, we used SNN Model$_D$ and Model$_E$. The population in both datasets was the same, yet the behavioural experiments they participated in; had slight variations affecting the reinforcement training arrangement (Table 20). Again, generalisation tests were run with the remaining four datasets. Overall, the performance ranged from 55%–74% in identifying the component. In particular, generalisation tests run for SNN Model$_D$ outperformed predictions made when using SNN Model$_E$ (Table 25).

In conclusion, the SNN Model$_A$ and Model$_D$ used more samples (Table 21), indicating a better exploration and exploitation of the search space, driving the algorithm to a better approximation of the global optima. This outcome was reached regardless of the population or the differences in the components. Further, patterns found in Dataset D can be hidden in Datasets A and B. On the other hand, Dataset E was derived from experiments with less similarity in the components than any different dataset in which variability contributed to the prediction ability to decay.

Table 25

Results of Weighted F1-score from Testing all *SNN* Baseline Models with Various Datasets

Case Studies	*Baseline Model*	*Datasets for testing*	*Mean*	*STD*	*Characteristics of datasets participating in training and testing the SNN models*
1-Datasets A and C were extracted by different populations, same RA	Model$_A$	*B*	*0.84*	*0.01*	*Same population of pigeons, same reward arrangements (RA)*
		C	*0.64*	*0.04*	*Different population, same RA*
		D	*0.72*	*0.01*	*Different populations, different RA*
		E	*0.75*	*0.02*	*Different populations, different RA*
	Model$_C$	A	*0.72*	*0.00*	*Different population, same RA*
		B	*0.66*	*0.01*	*Different population, same RA*
		D	*0.66*	*0.01*	*Same population, different RA*
		E	*0.69*	*0.01*	*Same population, different RA*
2-Datasets D and E were extracted by the same populations, different RA	Model$_D$	A	*0.75*	*0.06*	*Different populations, different RA*
		B	*0.74*	*0.06*	*Different populations, different RA*
		C	*0.59*	*0.04*	*Same population, different RA*
		E	*0.68*	*0.08*	*Same population, different RA*
	Model$_E$	A	*0.55*	*0.01*	*Different populations, different RA*
		B	*0.65*	*0.03*	*Different populations, different RA*
		C	*0.56*	*0.01*	*Same population, different RA*
		D	*0.68*	*0.02*	*Same population, different RA*

Key: Same RA= same procedure was applied for the pigeons participating in different behavioural experiments that occurred at different times; Different RA= different procedure each time, where reward arrangements parameters are dissimilar (see Table 20).

All datasets used to test for generalisation did not participate in the training of the tested model (for example, Dataset B was tested with the model trained with Dataset A denoted as Model$_A$ ▻B).

Further, confusion matrices were analysed across different populations under different components (Figure 43). The results here validate those obtained by (Plessas et al., 2022), where similar patterns can be hindered in close proximity reward arrangments represented by similar than when the components are presented in extremes. Yet, the prediction performances vary depending on the characteristics of each dataset used either for training or testing.

Figure 43

Example of confusion matrices of different generalisation tests with SNN Model$_A$ and Model$_C$ with the remaining Datasets. Model$_A$ is presented in the left column with the letter A and Model$_C$, right column, as C.

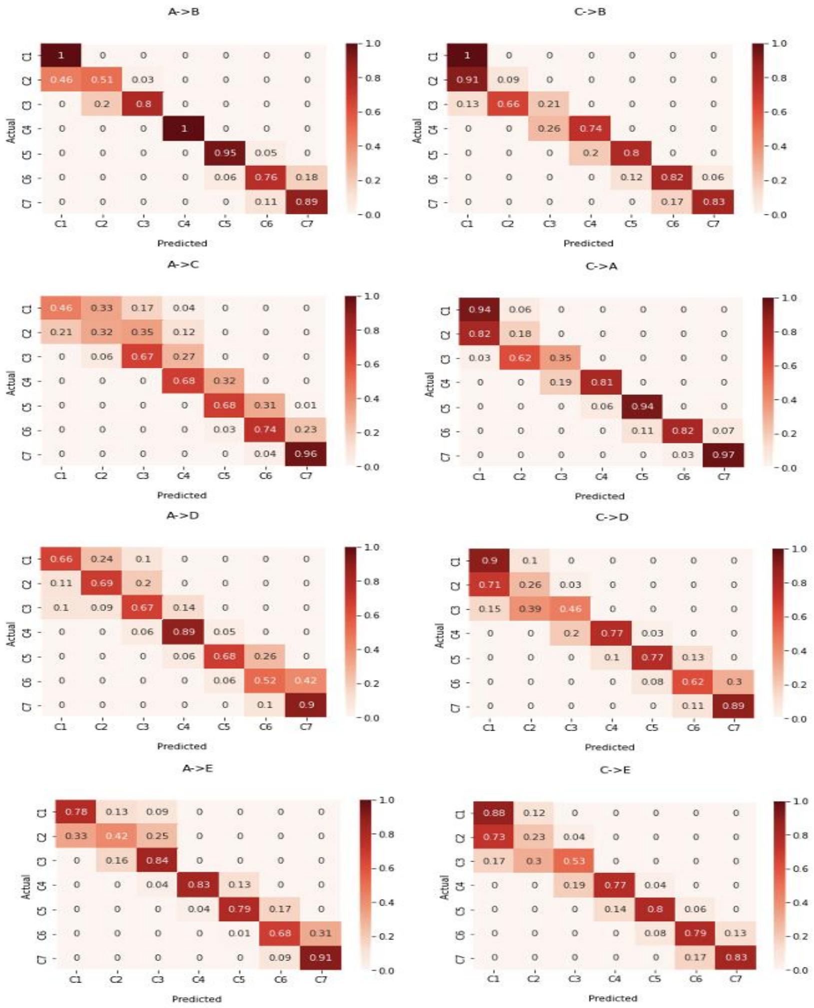

7.8.2 Comparison between transfer learning strategies for SNN models'

The performance of both transfer learning strategies was evaluated for all SNN models (Table 26). Four different models were tested, applying both transfer learning strategies (i.e., feature extraction and fine-tuning) when using five different datasets as inputs. As mentioned above, datasets were extracted by experiments using different animal populations and/or different reward arrangements. Both transfer learning strategies produced high predictive performance when tested on other datasets (Table 26). In particular, the lowest prediction performance was with SNN Model_E when predictions were made for Dataset C for both transfer learning strategies. Yet transfer learning fine-tuning demonstrated the highest level of performance   by achieving the highestlowest score (65%) with uniform distribution; an 11-point upgrade from the prediction ability compared to feature extraction for the same dataset. The highest results were with SNN Model_A and Model_C and Dataset B. By this, we conclude that taking the underlying patterns (weights and LIF parameters) and adjusting only the outputs (i.e., firing rates) can be sufficient to predict the component. When applying fine-tuning that required the output firing rates to be adjusted to the new firing rates to predict the different datasets, uniform slightly outperformed the other two approaches.

When comparing both transfer learning strategies, we observe fine-tuning slightly outperformed feature extraction. Yet, this doesn’t indicate that one strategy is better than the other, as both were able to produce a predictive outcome in close proximity. However, it’s worth mentioning that the time needed to train the baseline model with each strategy was sufficiently different, with transfer learning feature extraction producing the outcome in seconds as the outcome only required to be adjusted (i.e., re-calculated) to the new dataset. In contrast, with transfer learning fine-tuning, the baseline model had to be retrained by using the differential evolution strategy setting the number of iterations to 150.

Table 26

Results of Weighted F1-score for Both Transfer Learning Strategies

	Validation Tests	Feature extraction		Fine-tuning					
	Datasets			Uniform		Bounded		Normal	
		Mean	Std	Mean	Std	Mean	Std	Mean	Std
Model$_A$	B	0.84	0.01	0.86	0.06	0.85	0.10	0.83	0.07
Model$_A$	C	0.66	0.01	0.68	0.08	0.68	0.08	0.69	0.10
Model$_A$	D	0.71	0.01	0.72	0.08	0.72	0.10	0.66	0.07
Model$_A$	E	0.76	0.02	0.77	0.09	0.76	0.11	0.77	0.07
Model$_C$	A	0.83	0.00	0.83	0.05	0.81	0.07	0.82	0.06
Model$_C$	B	0.84	0.00	0.86	0.08	0.86	0.09	0.86	0.09
Model$_C$	D	0.72	0.07	0.72	0.05	0.71	0.11	0.72	0.08
Model$_C$	E	0.77	0.00	0.78	0.09	0.80	0.08	0.74	0.09
Model$_D$	A	0.80	0.05	0.82	0.08	0.81	0.04	0.83	0.06
Model$_D$	B	0.81	0.05	0.84	0.11	0.85	0.08	0.85	0.07
Model$_D$	C	0.60	0.06	0.68	0.08	0.68	0.07	0.67	0.07
Model$_D$	E	0.68	0.10	0.81	0.06	0.80	0.08	0.81	0.10
Model$_E$	A	0.61	0.02	0.83	0.06	0.83	0.05	0.65	0.07
Model$_E$	B	0.68	0.03	0.77	0.08	0.75	0.14	0.69	0.13
Model$_E$	C	0.53	0.01	0.65	0.11	0.59	0.12	0.56	0.06
Model$_E$	D	0.69	0.02	0.70	0.05	0.66	0.07	0.69	0.10

7.8.3 Comparison between baseline and target SNN models with transfer learning strategies.

We observe that the transfer learning feature extraction and fine-tuning approach in this study (Table 26) either produced the same or improved predictions (Figure 44) when comparing the prediction performance for each baseline SNN model, with each combination of datasets for all generalisation tests (Table 25).

Figure 44

Comparing predictive performances across pairs of baseline and target models used for transfer learning

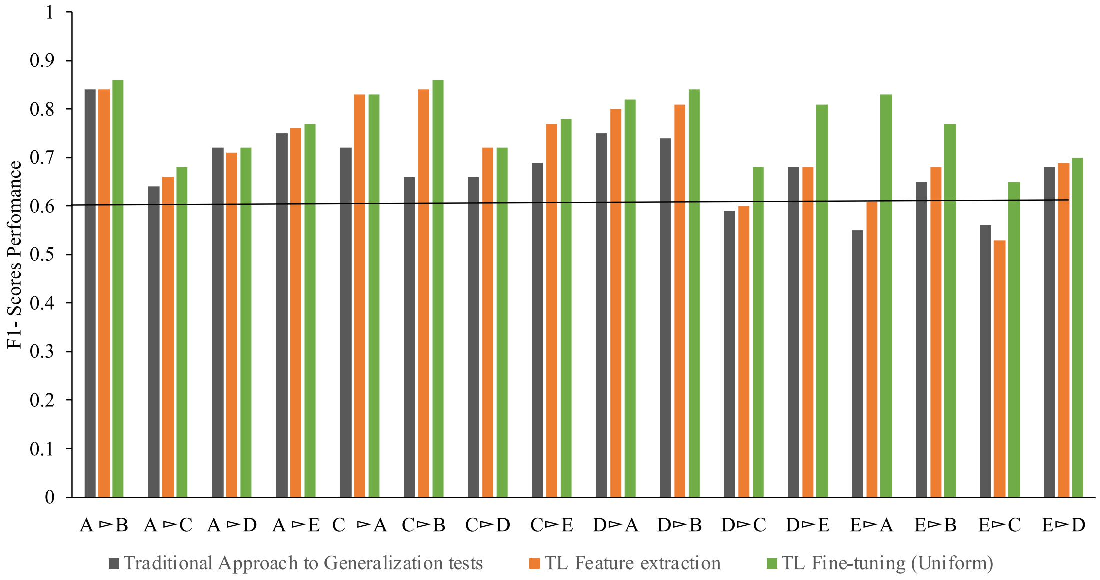

Further, we compared the time needed to train the model using nested cross-validation with transfer learning fine-tuning. The adaptation time to train the target model to generalise from different (or the same) animal populations and from different (or the same) reward arrangements was depreciable. Overall the predictive outcome was achieved faster with transfer learning fine-tuning requiring only half iterations for the training (Figure 45). No iterations were required with transfer learning feature extraction as no training was needed and the output was adapted to the new dataset. Also, with transfer learning, the algorithm had evolved with baseline SNN model training; therefore, any training for the target SNN model with 5-fold cross-validation to calculate the firing rates started at a higher point for training and testing. Hence, though prediction outcomes for traditional machine learning were in close proximity when applying transfer learning, it demonstrated that it could work even with less computational power. This can be important when using low-budget computational units such as mobile devices, when working on other datasets. In the example presented in Figure 45, one can view that learning via fine-tuning can improve with just a few iterations.

Figure 45

An example of the evolution of the differential evolution algorithm during the training of the SNN model either with nested CV and SNN Model$_A$ transfer learning fine-tuning when tested with four datasets (B, C, D, E).

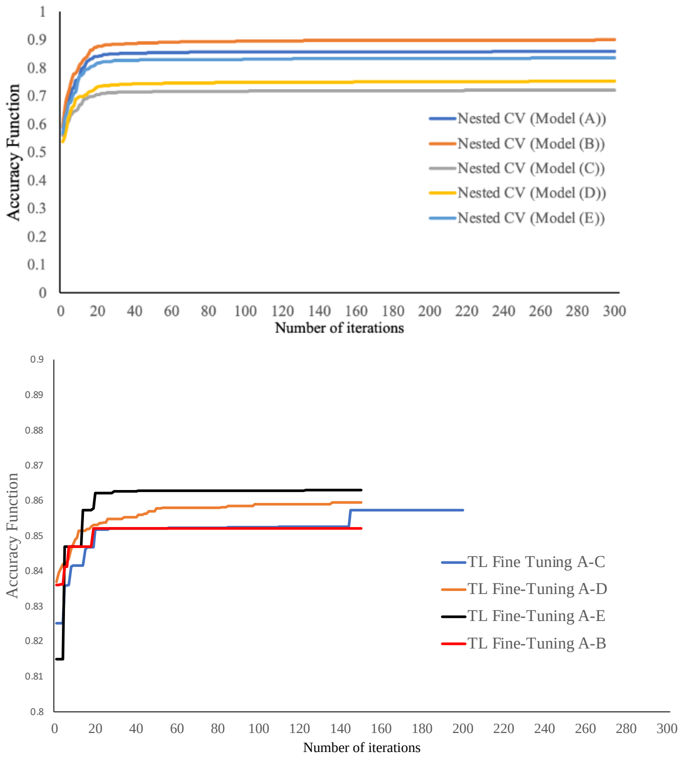

In conclusion, with both transfer learning strategies having good starting points they were able to expand the search space, resulting in a better approximation of the global optima. Thus, when using transfer learning strategies with SNN, the data allowed the model to become flexible and expand the search space if needed. It may be possible that models whose performances are evaluated solely on accuracy, can reach good predictions yet may be rigid. Therefore, transfer learning allows the models to become flexible with new datasets. For example, when the baseline model was developed with Dataset A, the transfer learning strategies allowed the target model to reach the best outcome with all datasets in just a few iterations across an animal population or experimental procedure.

7.9 Conclusion

Behavioural psychologists examine the functional relations between environment and behaviour by manipulating reward arrangements and observing changes in organisms' behaviour. Our study demonstrated that transfer learning methods could be suitable for studying animal behaviour. Further, our study showed that both transfer learning techniques (i.e., feature extraction and fine-tuning) could reach similar or better predictions than traditional machine learning modelling in identifying the training arrangement, by detecting patterns in the pigeon’s binary choices regardless of the dataset used. This outcome can be achieved rapidly, allowing predictions to be made on low-budget devices, thus, giving accessibility to behavioural researchers to use SNN models.

We conclude that using SNN models with transfer learning when dealing with small variable data (non-imaging and non-physiological data) can be critical. Furthermore, these models can become a valuable tool for studying environmental-behaviour interactions in behavioural analysis research, allowing the SNN model to tailor to specific organisms’ variability (‘data response model’), by demonstrating that the outcome is consistent across animal populations and the experimental procedure used. Therefore, the artificial model is adaptive, targeting the data of individual organisms and allowing the study of the individual organism, which is the goal of behavioural research. Also, adaptive modelling enables the model to become flexible, thus more biologically plausible. Future work should look into other datasets and other species.

Chapter 8 Overall discussion

The present experiments investigated whether a spiking neural network could learn from small samples of binary choices to predict and more precisely, retrodict learning histories unknown to the artificial model. By introducing a new way to use behavioural responses alone, this (non-traditional) approach is novel and innovative and demonstrates the opportunity to rethink the attainability of finding answers to questions that were not posed hitherto or did not get traction. The integration of SNN predictions to solve existing problems has the potential to offer new insights and applications for research, therapy design and treatment.

In Plessas et al. (2022, Chapter 4, p.69), the goal was to assess if an SNN could be trained to detect patterns from two-alternative choice responses (input data) extracted from a 2-key concurrent schedule procedure to make retrodictions about the learning history. An SNN was chosen due to its distinct ability to process temporal discrete events using spikes, rather than treating the data as a continuous accumulation of responses like other ANN models. By incorporating time dependency and leveraging spikes, the SNN was deemed suitable for identifying learning histories. The aim was to ascertain whether learning histories could be identified from discrete events constituted by a small window of 5 seconds of choice behaviours. In addition, we evaluated the model's accuracy on a new dataset comprising binary choices from the same subjects to assess how well it could handle different small-sized datasets. Experiment 2 (Chapter 5, p.109) aimed to investigate the SNN performance compared to other ANNs. The performance of the RNNs and CNN was compared to the trained SNN in predicting (or retrodicting) the unknown histories that had shaped their responses. This step was necessary to validate the efficacy and novelty of the SNN approach we adopted in Plessas et al. (2022, Chapter 4, p.69). Experiment 3 (Chapter

6, p.141) expanded the investigation by demonstrating that an artificial model's prediction can aid in identifying the unique contingencies experienced by individuals, based on the detailed patterns of their learning history. The baseline model's network parameters (including weight threshold and refractory period) were applied to modify the spiking neuron's prediction when presented with unique patterns from new datasets, to achieve personalised modelling with each new dataset. The approach represented an advancement that can account for generalisation from different participants and learning histories. The last Experiment (Chapter 7, p.175) aimed to retrain the SNN model using all datasets and applied two transfer learning strategies to assess the SNN's generalisation ability. These experiments delivered a trustworthy single-neuron SNN that could be used easily and efficiently, opening new avenues for behavioural research and clinical practice.

(Table 2 presents a summary of all research studies.)

8.1 SNN for behavioural research and clinical practice

The SNN, like any artificial neural network model, makes predictions based on the patterns and the features of the data it was trained on. The fact that the model developed in this thesis made accurate predictions implies that it could handle binary choices. The performance metrics, including accuracy, recall, and specificity, shown in Table 4 of Plessas et al. (2022, Chapter 4, p 69), indicated that the SNN model performed effectively by achieving scores above 80% for all true positive (learning histories) and negative events (non-histories). The suitability of the SNN model was further strengthened by comparing its prediction performance with that of other deep learning models (CNN, LSTM, GRU), with the specific configuration when using the F1 measure. The results showed that the SNN outperformed the other models, demonstrating its effectiveness in predicting learning histories (Figure 24). If we relied solely on accuracy scores to evaluate model performance, we would be assuming that all models are equally proficient in handling the data. However, this would overlook crucial data characteristics (Kuhn & Johnson, 2018) that are essential for completing the task at hand. For example, we would have missed that all DNN models possibly had a bias towards one of the learning histories (Component 1), resulting in high accuracy for that learning history (Table 10) but a low F1 score. A low F1 score indicates that the model has many false positives or false negatives, which can affect the model's overall performance.

Using an SNN to model choice behaviour datasets, has several advantages over other ANN models. An SNN prediction that conforms to the generalised matching law (Baum, 1974) would be expected to show high recall and specificity measures, in addition to accuracy, when modelling choice behaviour datasets. This is because recall and specificity reflect the model's ability to correctly identify positive and negative cases, which is important for matching the proportion of responses, to the proportion of reinforcers obtained from each alternative. The SNN model reached high metrics on both values (recall≥81%, specificity≥96% (see Table 7), while the deep learning models underperformed with recall (recall (CNN) 48%-65%, (LSTM) 66%-76%, (GRU) 53%-78% (see Table 11-Table 13 in Chapter 5). This suggests that the deep learning models could not capture the data's essential characteristics when processing inputs as continuous values, and further investigation and optimisation would be required to improve their performance. By contrast, SNNs, at least in this project, did not require further training, making them a simpler and more efficient choice for identifying learning histories. Another advantage of SNN with choice behaviour can be the inherent ability to process time in the spiking activity (e.g., Gerstner & Kistler, 2002; Maass &

Markram, 2004; Tavanaei et al., 2019) . As the SNN was trained solely with samples of left- and right- choice behaviour that occurred in time, it learnt from information that had current value every time a choice was made. This is because the SNN is designed to process information over time, as the neuron fires at specific time points based on the input it receives (e.g., Rafi, 2021). Research on choice behaviour has demonstrated that reinforcers are stimuli/events in the organism's learning history that have been accumulatively correlated with events that follow them. Therefore, reinforcers signal future likely events (Cowie et al., 2013, 2016; Cowie & Davison, 2016a). This temporal relationship between behaviour and learning history is captured by SNNs. If all behaviours are choice behaviours, as proposed by Herrnstein (Herrnstein, 1961, 1970), this suggests that the SNN's ability to capture the spatiotemporal nature of data effectively, may have broader applications across other behavioural paradigms. In contrast, artificial models that use continuous values of data to train the model, such as RNN and CNN, may be better suited for other types of behavioural datasets that rely on data from behaviours occurring within a limited timeframe (e.g., datasets extracted from discrete trials procedures). Thus, predictive models, like RNN, can utilise inputs generated from a series of behavioural occurrences near a reinforcer without the need to consider information beyond a limited timeframe that affects the behaviour in question. This viewpoint aligns with Skinner's temporal perspective (Ferster & Skinner, 1957; Skinner, 1938, 1969). For example, for a functional behaviour analysis at a given moment, one can use data of the immediate contingencies manipulated to investigate what controls the behaviour (Iwata et al., 1994). However, if learning histories are taken into account, algorithms that capture the temporal dynamics of the behaviour-reinforcer relationship would require information from patterns of behaviour-reinforcer interactions over time (Cowie, 2018; Cowie & Davison, 2016a) by including distant learning events. Therefore, datasets including prior, historical events would be better processed by SNN.

Nevertheless, RNNs have been reported in the literature for handling temporal variables (e.g., DiPietro & Hager, 2020). As Experiment 2 (Chapter 5, p.109) demonstrated, additional training may also be valuable for behaviour analysis. On the other hand, CNN showed some overfitting with binary choice behaviour and predicting learning histories, which could relate to the fact that CNN is better at handling images. Therefore, the spatial input of the data seems more critical for the algorithm to learn, as they are designed to handle 2D spatial relationships (Li et al., 2022). It is possible that in choice behaviour datasets, temporal information is more important for making predictions than the specific location of the response (left or right-key). Therefore, CNN may become an option for images or videos of binary behaviours that constitutes the inputs (i.e., grid-type data). In conclusion, with no need for additional training, SNN seemed more potent with binary choices than conventional ANNs as they are time-defined.

Moreover, it was evident that the SNN effectively handled small datasets of pigeon subjects with different learning histories, containing variable discrete events. As the number of reinforcers changed in the experimental conditions, the time points, that is the points of post-reinforcer 5-sec periods, also varied, and therefore, the sample size varied too. However, the number of samples generated was approximately the same (Table 3, Table 24). In Plessas et al. (2022, Chapter 4, p.69), when generalisation was tested using conventional methods (Chapter 6) on a dataset that was presented with three different sizes (45-, 63-, 90- periods of 5-s), the prediction was persistently high across all metrics (Table 8). The same was true when datasets were used for different pigeon subjects and when the learning histories had variations in the probabilities of receiving a reinforcer with a concurrent schedule procedure (Table 19, Table C3-C4 Appendix C, p.301-302).

However, the same did not apply to the dataset extracted from Condition 1 by Davison and Baum (2000) (Table 15). In particular, recall values were affected the most (40%-57%). It would be hard to assume that the SNN model may have overfitted to the datasets as the model could still generalise to novel small datasets with variant characteristics (Condition 4 and 9, Table 19). Alternatively, it is possible that differences in the distribution of response patterns in Condition 1 were different as presented across time and, thus, had a wider range of values or patterns when compared to the other datasets. Therefore, one interpretation may be that the dataset itself had more variations in patterns, that were not easily detected. In the source study (Davison & Baum, 2000), they found that the effect of reinforcers upon choice behaviour changes as the distribution of within-session changes in components occurred. The datasets were analysed, and Table 15 revealed a discrepancy in the distribution of response patterns between the dataset extracted from subjects other than those used to train the model (Table 4); though the source from which the data were extracted, the studies concluded that long-term aggregations of reinforcers remain important, as do individual reinforcers and sequences of reinforcers.

A logical conjecture may be that the SNN model successfully handled small datasets, but its ability to generalise to small datasets with varying distribution of responses might hinge on behavioural response patterns either related or unrelated to the ongoing reinforcement contingencies. In saying that, generalisation can also be affected by the distribution of pattern shifts across species or by other behaviours. An initial experimentation with rat and human data extracted from other studies (Krägeloh et al., 2010; Landon et al., 2007) and from experimental conditions identical to the Condition 1 dataset, revealed that the personalised modelling approach could account for differences in behavioural responses represented in the data (Appendix C). This was evident when the model successfully extended its predictions from pigeon binary-choice datasets to rat binary-choice datasets (with the rats choosing between left and right lever) (Landon et al., 2007). Similarly, another experiment extended the SNN model’s predictions from pigeon binary-choice datasets to human datasets, where the human subjects were choosing between the left and right side of a screen to pop balloons (Krägeloh et al., 2010). The SNN model's ability to transcend species and interaction methods underscores its robustness and potential applicability across diverse contexts. Furthermore, mathematical models have demonstrated that the underlying mechanisms of choice behaviour are conserved across different organisms, including mammals (Cowie et al., 2021; Podlesnik et al., 2022). Therefore, it is reasonable that the SNN personalised approach is generalisable to small datasets either within the same or other species, potentially informing us of ontogenetic and phylogenetic contributions. It is possible that variations in predictions (observed when training the artificial model on pigeon datasets and testing it on other species) can potentially unveil distinctions between ontogenetic and phylogenetic behaviours. This would provide valuable insights into how ecological factors influence learning processes. For example, the predictions obtained by the SNN (Appendix C, Tables C3) might suggest that, despite differences in anatomy and evolutionary history, there are behavioural contexts where pigeons, rats and humans exhibit similar patterns of learning, in line with the view that species have some shared mechanisms of learning. However, further research and validation are required to either confirm this or provide support for the proposition of organisms ‘misbehaving’ due to phylogenetic influence (instinctive behaviours) (Breland & Breland, 1961). In conclusion, when an SNN model generates less accurate predictions –as observed with the dataset extracted from Condition 1 (p.152)– it may strongly reflect the characteristics of the dataset itself, which might have emerged either from phylogenetic predispositions or from ontogenetic biases.

The usual practice dictates that if a simple model did not capture the data well, predictions would drop, and one should consider using a more complex SNN model to improve performance. However, increasing model complexity may not be cost-effective. Also, to avoid the model overfitting with the data, additional training data are usually required to create large datasets (Alpaydin, 2014, 2020). In contrast, smaller datasets offer greater opportunities for closely monitoring (Orrù et al., 2020) and capturing the heterogeneity in choice behaviour. As they provide insights into the complexity of the data, such as choice behaviour (Cowie, 2018, 2019), these findings offer a perspective that supports scaling down, with small being supported from other disciplines such as psychiatry (Koppe et al., 2021) and clinical predictions (Riley et al., 2020).

We adopted a unique strategy for behaviour analysis, by preserving the model's simplicity and leveraging transfer learning, while also continuing to work with limited datasets rather than following the trend of other behaviour analysts, as well as psychology in general (Adjerid & Kelley, 2018), to switch to big datasets (e.g., Cox et al., 2021). Transfer learning involves making use of a pre-trained model on a dataset to improve predictions on a new model (Weiss et al., 2016), particularly when training and generalising to a dataset with different distributions than the original data (Niu et al., 2021). Thus, the model exhibited the ability to learn and accurately solve a similar problem, and subsequently utilised this knowledge to adapt faster to a new problem by capitalising on the lessons learnt from the initial problem. When applied transfer learning methods were used, in Experiment 3 (Chapter 6, p.141), Experiment 4 (Chapter 7, p.177), regardless of the method of the SNN training (cross-validation or nested cross-validation), the SNN improved in its ability to capture the complex relationship between input and output. Consequently, the SNN adjusted to better account for the variability in the input samples. For example, improvements were observed in Condition 1, originally described by Davison and Baum (2000), with the highest data variations. Furthermore, transfer learning when using Condition 1 datasets, led to better learning history predictions and reduced the number of false positives and negatives in both Experiment 3 Chapter 6 (F1 scores increased from 0.40 to 0.67) and Experiment 4 Chapter 7 (F1 scores increased from 0.64 to 0.69). Nevertheless, it is possible that no further improvement in predictions with this dataset could be attributed to the fact that patterns may not have been detected within the small observation window (5-sec after reinforcer delivery) or other factors related to the SNN model.

We utilised a personalised approach to assess inter-subject variability, testing individual datasets. For example, Table 18 demonstrated these differences with datasets extracted from Condition 1 from Davison and Baum (2000). This allowed us to observe that the approach remained effective, as diverse predictions can be attributed to individual differences in learning histories and their contribution to the overall group prediction. Individual predictions may be accurate and informative, but when combined as a group, they may not be as effective, due to the variability among individuals. In summary, achieving good generalisation results is a key goal in machine learning, and it remains a challenging task that requires careful consideration of the data, model, and training process (Alpaydin, 2020). Therefore, it is demonstrated that when working with different sizes of small datasets that may show variance, it is important to consider both individual and group predictions in any analysis. Our approach showcased the SNN's ability to remain simple and flexible without needing more complex models. This highlights the model's adaptability to capturing spiking activity represented by its firing rates. Future work can investigate more parameters to capture the patterns in the data, which are discussed later.

8.2 Theoretical contribution (Reinforcer-behaviour relations)

When experimental procedures involve frequent changes to environmental parameters such as reinforcer ratios, overall reinforcer distribution, and the number of reinforcers, the behaviour can change quickly. Behaviour learnt can be affected by elapsed time and where it happens, provided that the environmental change becomes highly relevant to the animal's behaviour (Davison & Cowie, 2022). This thesis is one of the few examples of using an artificial model to gain insight into this behavioural paradigm (Figure 3). Until now, this project is the only one that has demonstrated the potential of using spiking neural networks to analyse choice behaviour data. Furthermore, adopting a unique approach that treats learning histories as unknown and factoring in the prediction time elapsed and location of response, gave valuable insights into the relationship between reinforcer and behaviour, as discussed in the following section.

Measures of accuracy are commonly used to evaluate the performance of artificial neural networks (ANNs), including spiking neural networks (SNNs). However, some researchers have explored the inclusion of more measures to better understand how well the network is performing (e.g., Jordan & Mitchell, 2015). In all Experiments (Chapters 4-7), the performance of an SNN was evaluated using a similar approach, going beyond accuracy to assess the performance with additional measures. As a result of this more comprehensive evaluation, the SNN's performance was satisfactory on these additional measures (for example, see Figure 23 when DNN models were compared with accuracy and Figure 24 when relying on F1 scores). Therefore, in Experiment 3 (Chapter 6, p.141), I demonstrated that precision values directly impact the trustworthiness of the SNN model in retrodicting learning histories and can be beneficial when combined with recall and F1 values. Consequently, being able to predict (retrodict) the learning history 'truly' reconfirmed the reality of the reinforcer-behaviour relationship regardless of the theoretical explanation for this relationship (e.g., Simon et al., 2020 for the current debate on the mechanisms of this relation).

Moreover, behaviour analysts have been using mathematical models to extract and analyse data for a long time. However, the use of artificial intelligence models, such as spiking neural networks (SNNs), has the potential to take this analysis to the next level. Because SNNs can extract and analyse information automatically, they are capable of learning to associate environmental conditions with behavioural patterns forming their own stimulus-behaviour-reinforcer relationship (Alpaydin, 2020). This raises the question of whether SNNs could learn the same way as a pigeon, by the model making decisions based on binary choices, that inform us about the environment (e.g., Cyr et al., 2014). The ability of a simple SNN model to adapt to the data caused by the environmental changes, was demonstrated by the use of new input samples and different patterns in time. As seen in Chapters 6 and 7 with pre-trained models and datasets, it means that it has the capacity to make 'computational decisions' based on the data and the experience gained from previous learning and the current datasets. The consistently high performance of this SNN when using binary choices suggests that choice behaviour 'carries' information about past learning, and also, the nature of the data allowed the spiking activity to learn. This finding creates potential for future research in this area.

If we are to utilise SNN to understand how adaptive learning takes place and choice behaviour patterns are shaped, we need to use the predictions generated by the model as a starting point to develop hypotheses and investigate how different factors interact during learning. The predictions made by the SNN from individuals' datasets (Table 18) and groups of different subjects and the variations in their learning histories (all experiments) lead to the question of whether predictions based solely on snapshots of binary choices can provide any additional information about the relationship choice behaviours have with the reinforcer in these frequently changing procedures. Beyond that the relationship is quite apparent. Thus, differences in the level of success the model achieved during prediction-making, raised questions about explaining those differences. A closer analysis of the metrics and linear regression data in Plessas et al. (2022, Chapter 4, p.69) Chapter 4 reveals that the firing activity of the neuron in question may not have been entirely driven by the relationship between components and choices made. As shown in Appendix B (B4, p.295), the values of sensitivity (a) to reinforcement from each bird and the inherent bias (log c) could have facilitated either a better prediction (e.g., Component 7) or the contrary (e.g., Component 2). What is illustrated is that individual pigeon differences in patterns of responding, seem to play a role with the model's retrospective ability. Possibly, the increase in sensitivity to reinforcement by itself may not ensure better predictions (retrodictions) when influences from bias are present. However, the less sensitivity pigeons showed to reinforcement, the harder it was for the SNN model to identify patterns and make predictions.

In our analysis (Appendix B), the sensitivity values started high, which can be explained by the small window of observation. Pigeons could be reusing the just-reinforced response when they signal future reinforcers (Gomes-Ng et al., 2017; McLean et al., 2014) and the SNN model’s retrodictions suggesting something about choice behaviour (Chapter 6). Moreover, between-pigeon variability is high, indicating that each pigeon responded differently to reinforcement, which also reflects on the variability of the machine's metrics in the testing phase (Figure 13). The fact that the model detected a pattern and identified component-related choice behaviours indicates that the pigeons' responses are not just 'reinforcer pulses', i.e., repetitions of a just-reinforced response. The temporal proximity to the just-reinforced response should have a significant effect on choice, thus controlling subsequent responses (e.g, Davison & Baum, 2003; Krägeloh et al., 2005; Landon et al., 2003a, 2003b). However, if they were simply copying the just-reinforced response, there would be no inherent bias (see GML), their responses would strictly match reinforcement, and components would be easy to discriminate. This argument is also strengthened by the results obtained in Experiment 3 (Chapter 6, p.141); precision scores would expect to be higher if the subjects were copying the most recent reinforced choice behaviour. The analysis of the predictions suggests that the timing of reinforcer deliveries in time and the location the response occurred (left or right) could influence the type of information the subject has about where future reinforcers were likely to occur. This hypothesis aligns with current research on choice behaviour that has demonstrated time's effects on pigeons' behaviour provided that the subject discriminates which response provided the reinforcement (e.g., Cowie et al., 2011, 2013). Thus, detecting the temporal relation between reinforcer-behaviour in time may be a critical factor for predictive modelling of learning histories, as research implies (e.g., Cowie et al., 2014; Davison & Cowie, 2022).

Spiking neurons transform data into spikes over a specified temporal window. With time being an inherent feature of spiking activity, it is possible that the output of the single-neuron SNN reflected the timing of reinforcer deliveries. As every time point in our samples represents choice behaviour immediately after the reinforcer delivery, and since predictions by SNN were variable, that indicates choice behaviour was not always favouring the just reinforced behaviour. So, changes in choice behaviour patterns may reflect whether the subject was at least partly able to discriminate how accumulated reinforcers work overtime, not just at the moment. Also, this ability 'to sort out' how the reinforcer works may be possible even in environments like the ones used here, where no specific signal was arranged for when the reinforcer was delivered and changes within sessions were frequent. For example, in Experiment 3 (Chapter 6, p. 141), choice behaviour datasets extracted from Condition 4 denoted less data as fewer reinforcers were provided, whereas, in Condition 9, the rate of reinforcer deliveries increased (more time points) (Table 15). The result was a slight increase in the precision of the prediction when speed increased in that environment (Condition 9 dataset).

On the other hand, predictions were low in the dataset extracted from Condition 1 despite the increase in the data in each sample (nine reinforcer deliveries). These results may indicate that the number of reinforcers did not affect the model's predictability. Combining the results from all experiments, seems that different factors in each pigeon's learning history contributed to the SNN's ability to retrodict their choice behaviours. The values of sensitivity the pigeon showed to the reinforcement, the overall reinforcer ratio it was exposed to across time and its individual preference unrelated to reinforcement each played a role in retrodictions. These findings, though, could suggest that the effect of a reinforcer on behaviour may be better understood by how discriminable the reinforcer is and how reinforcers are given over time and space, which in turn affect how the organism responds to the environment. These characteristics in the datasets were identified by the SNN model explicitly based on left and right responses and without feeding the SNN information about reinforcers were delivered.

As the datasets used with the SNN included a small window of choice behaviours, the data points represent choice behaviours made around a reinforcer which partly made the reinforcer a discernible marker carrying more information than initially thought, when arranged by different reinforcer arrangements. Furthermore, since the SNN reflects the passage of time in its spiking activity, spatiotemporal correlations were partly detected in behaviour patterns when responding to post-reinforcer deliveries. Therefore, differences in the SNN prediction may be supported by the literature: if the subject is exposed to reinforcers over time, then the SNN may be discriminating the 'discriminability' – in other words the animal's learning while in training that differences in the contingencies taking place– and working out which choice behaviour provided the reinforcer with respect to the elapsed time. If this task becomes challenging, the subject can be unsure which behaviour produced the reinforcer (Cowie et al., 2014; Davison & Jenkins, 1985). The SNN's variations in predictions could stem from the subject's ability to discern which behaviour led to the reinforcer over time, indicating that the SNN mimics the pigeon's behaviour. Since SNNs draw inspiration from biology, exploring the connection between learning and behaviour might lead to neuron adaptation and increased flexibility. This adaptability could involve incorporating spatiotemporal variables to capture the bird's ability to discriminate between contingencies.

In summary, the research presented in this thesis, demonstrates the usefulness of predictions (retrodictions) made by a dynamic artificial model (SNN) in studying how reinforcers influence behaviour and contribute to predictions of learning histories. Specifically, a Leaky Integrated-and-Fire (LIF) model (Appendix B, p.285), which is a deterministic model that accounts for response-to-response variability (Gerstner & Kistler, 2002), can generate *informative* predictions that can point to new hypotheses used in future research. By analysing the spatiotemporal information contained in choice behaviours made by subjects in response to the left- and right-key arrangements, a snapshot of data revealed the potential role of time elapsed and reinforcer discriminability, when learning histories are approached from a different perspective. These findings create exciting opportunities for future work and the development of SNN models.

8.3 Clinical contribution

The ability of the SNN model to retrodict learning histories showed that limited observational behaviour data, which is the usual case in naturalistic settings, allows the clinician to shape future behaviours without requiring a detailed understanding of how the functional relations were established in the environment. The application of AI assists clinical practice allowing EAB research to become usable for specific purposes.

The use of AI in EAB research has the potential to extend beyond the development of theories and impact the lives of those in need. However, the complexity of EAB research may pose a challenge for clinicians to implement its principles effectively. It has also been long reported that identifying the environmental condition that participates in behaviour change is a challenge in therapy (Lattal, 1995; Thompson & Iwata, 2007). This thesis demonstrates the potential of SNN models to leverage the clinician's competency when arranging reinforcer contingencies to create an optimal learning environment that is task specific and/or learner specific. Also, this thesis demonstrates the potential of SNN models to provide individual and group-specific information that can enhance decision-making in arranging reinforcer contingencies, even with limited observational behaviour data. The predictions generated by the SNN model can be used to identify which environment can effectively build future behaviour based on their learning history, ultimately improving the effectiveness in clinical practice.

Also, considering the spatiotemporal parameters, an SNN prediction could assist the clinician's decision-making process to be more accurate. This is because the predictions made by the SNN are based on 'dynamic choice behaviours, influenced by variable factors such as time elapsed in context with other stimuli available, which is not easily captured when designing a treatment plan. The Behaviour Analyst Certification Board sets the requirements for 'competent practice' and requires the behavioural analyst to plan and make decisions about the effectiveness of the treatment based on data (item, H, 5th Task List, BACB - Behavior Analyst Certification Board, n.d.). However, in 'real' practice, the treatment plan is set in advance, and any revision is based on the past (treatment) plan and the data collection acquired based on the plan. This creates a gap between the contingencies taking place when the plan was developed and those taking place as the plan is put into practice. An SNN tool that can predict learning histories in real-time during therapy, could potentially fill this gap by enabling more timely and accurate decisions, considering new behaviour inputs generated. However, further investigation is necessary, and future work will address this possibility.

8.4 SNN contribution

While this study was not focused on SNN engineering, it did make several contributions to the field. First, we demonstrated that simple behaviours, though more complex than they appear, can train SNNs. This is a novel demonstration, as the model used was developed and tested with static data (Vazquez & Cachón, 2010). In particular, the input patterns used for pattern recognition did not change over time. The patterns were transformed into input signals that were presented to the LIF neuron, which was stimulated for a fixed duration of time in milliseconds, and then the firing rate was computed. The synaptic weights of the neuron model were then adjusted based on these firing rates. In Plessas et al. (2022, Chapter 4, p.69) the SNN model was presented to learn, using dynamic data, which would involve analysing changing patterns or signals over time.

Further, when the outcomes were compared to other DNNs (Experiment 2 Chapter 5, p.109), it demonstrated how different neural networks work with behavioural observation data and possible limitations. This was novel, as such a comparison has not been demonstrated in the literature and there are limited applications of these forms of data (Dehghan et al., 2022). Lastly, applications of transfer learning techniques with SNN are rare in the literature and Experiments 3 (Chapter 6, p.141) and 4, (Chapter 7, p.175) demonstrated the applicability of SNN. Thus, these experiments have provided valuable insights into SNN methods overall.

8.5 Limitations and future prospects

While it is important to acknowledge that the results of this study may be limited by the specificity of the single-neuron SNN model and the datasets used, it is worth noting that other behavioural paradigms using similar procedures have similar consistent results from behavioural observations. Therefore, it is reasonable to speculate that the application of the SNN model could potentially have broader implications. Future work would require the SNN model to be tested further, considering the prediction’s utility. For example, if the aim is to use the SNN model in experimental research and use predictions to investigate the underlying spatiotemporal dynamics between reinforcer-behaviour, the SNN model can be tested with other datasets taken from other procedures (for example, steady state procedures, experiments manipulating time, high- and low-discriminability histories, or with datasets from the same subject across different experiments).

In this thesis, taking a further step and allowing the experiments to re-calculate the firing rates of the SNN helped to improve its generalisation to new datasets. However, it is not the only factor that affects generalisation. Firing rates refer to the frequency at which the neurons in the network generate spikes. They can be affected by several factors, such as the input signals, the network architecture, and the synaptic weights. If new datasets are applied, one wants to investigate further generalisation with new datasets. One possible way is by working on the firing rates, such as type of time encoding. For example, one could experiment with different network architectures and parameters that promote a more balanced and diverse firing across the neurons. Also, more features may be included as input (e.g., number of reinforcers). By building on this work, one can explore learning patterns detected in choice behaviour when learning histories are approached from different SNNs.

A second future direction for research is investigating whether SNN models can be used to develop learning markers. Despite binary choices appearing to be simple inputs, the experimental setup can make this type of task relatively complex when environments are frequently changed. As behaviour change involves learning in specific environments, it may be possible to use the predictions made by an SNN model, to develop learning markers that provide insights into the brain processes underlying learning and behaviour change in different environments. This study was limited in that it only investigated 5-sec periods of choice behaviours immediately after the reinforcer. It may be possible that different time periods affect predictions differently. The next step would be to assess prediction performance using data from different time periods to identify the most informative time period for making predictions and inform us on critical learning patterns or times in learning. Moreover, predictive outcomes can inform the development of new hypotheses by detecting patterns that are also helpful in understanding a living organism's behaviour. In the future, treatments may provide more informative insights into the brain and learning patterns rather than solely relying on diagnosis to classify behaviours. For example, inaccurate timing of choice behaviour may result from pattern differences associated with neuropsychological conditions. This highlights the potential for treatment outcomes to serve as a valuable tool for understanding the underlying mechanisms of behaviour and learning, particularly in relation to how the environment influences changes in behaviour.

A third interesting future work would be to investigate further the model's data-responsive nature. A real-world application of SNN models would require the artificial model to learn continuously from small variable datasets occurring in real-time. The predictions can inform the clinician of changes in behaviour that occur in real-time relative to changes presently in the environment and based on past experience. In this thesis, the analysis was limited to assessing the performance of the predictions. For learning histories to find clinical applications to improve therapy (e.g., ABA interventions), further development is required. Inspired by the learning I achieved during this doctoral and the work conducted by Cyr et al. (2014), the motivation is to evolve this work. Cyr et al. (2014) demonstrated that complex computational models could be avoided, and they adopted a simple SNN model with few SNN neurons (synaptic links and a spike-timing dependent plasticity as learning rules) for the model to learn from reinforcers in multiple learning scenarios and when integrated into a robot, to have its behaviour change based on the environmental changes given.

A fourth area of application is more intensive use of tracking technologies. If we want the SNN model to actively change predictions based on behaviours that change in real time, an experimental paradigm can be developed to expand the SNN architecture to accept inputs that can be extracted from behaviours in real time. Using motor sensors for example could provide the clinician in actual time how to supply the reinforcers during a session and thus, affecting the behaviour based on a subject's personal learning history. An AI tool as such during therapy would leverage learning outcomes as the clinician, who is part of the environment, will also be adjusting reinforcers based on real-time changes in the client's behaviour. This makes interactions between reinforcer delivery by the therapist and behaviour change by the client having a feedback function, as there is constant interaction between the organism (client) and environment (therapist manipulations), whereby environmental changes are followed by the organism responding to them, thereby changing the environment, and so on.

Lastly, SNN have potential for theoretical-foundational research. Datasets from mammals were preliminarily tested in this thesis (Appendix C, p.300). They can be studied further to strengthen the performance of the model, and also to seek for answers to the question of control by ontogeny versus control by phylogeny. The utilisation of SNN has demonstrated promising results in emulating brain-like processing within specialised hardware, although their current capabilities fall short of replicating the complexity of the human brain. A comprehensive exploration necessitates experiments involving identical and diverse species datasets to understand an organism's responses to varied stimuli and corresponding behaviours. This approach enables the discrimination between behaviours linked to reinforcement (ontogenetic) and those genetically programmed behaviours (phylogenetic) that are not learnt through individual experience (i.e., instincts). Within this doctorate, the preliminary observations revealed its competence in delineating shared mechanisms of learning when the SNN model was trained on one species (i.e., pigeons) and tested on another (i.e., rats and humans). These findings could potentially guide us towards investigating further the variations in responses within a given species or the classification of discrepancies based on species-specific behaviours.

However, as the level of environmental control associated with the datasets becomes less stringent (e.g., with humans), identifying behavioural patterns through an artificial neural network may progressively escalate in complexity. Hence, a potential approach to address this issue involves implementing an elimination process to systematically focus on how many variables exert influence within this complex environment. Iteratively removing variables that introduce noise or ambiguity into the dataset might generate outputs that highlight phylogenetic versus ontogenetic patterns in the dataset. Consequently, the elimination process might enable the artificial model to focus on the core elements contributing to the ontogeny or phylogeny, thereby facilitating more accurate and meaningful analyses.

The SNN offers a data-driven perspective potentially providing insights without introducing prior theoretical assumptions, contributing to an objective and nuanced understanding of the underlying phenomena. Therefore, as SNN reflects patterns in data, one would expect it to incorporate both phylogenetic and ontogenetic behaviours. However, it is necessary to expand the scope of the experiments for a comprehensive assessment. The synergistic interdisciplinary collaboration among machine learning, behavioural science, and biology could help further advance knowledge in adaptive behaviours of various species in their respective environments that have developed through the interaction of environmental contingencies and evolutionary history.

8.6 Conclusions

Our experiments showed that a simple SNN model can detect patterns in small and variable data snapshots of current learnt behaviour and, importantly, when the data is extracted from unpredictable environments. This finding is significant because it demonstrates a different approach to analysing behavioural data by leveraging learning histories to support treatment and interventions and enhance theoretical understanding. By being able to predict learning histories, SNN models have the potential to revolutionise the field of behavioural science by providing valuable insights into the underlying mechanisms of learning and behaviour change. This information can then be used to develop more effective interventions and treatments that are tailored to each individual's unique learning history. Overall, these experimentations have added valuable insights to the field of SNN methods and have the potential to shape the future of behavioural science research.

References

4.1 Integrate-and-fire model. (n.d.). Retrieved May 1, 2023, from https://lcnwww.epfl.ch/gerstner/SPNM/node26.html

Adjerid, I., & Kelley, K. (2018). Big data in psychology: A framework for research advancement. *American Psychologist*, *73*(7), 899–917. https://doi.org/10.1037/amp0000190

Ahmad, M. F., Isa, N. A. M., Lim, W. H., & Ang, K. M. (2022). Differential evolution: A recent review based on state-of-the-art works. *Alexandria Engineering Journal*, *61*(5), 3831–3872. https://doi.org/10.1016/J.AEJ.2021.09.013

Alaloul, W. S., & Qureshi, A. H. (2020). Data Processing Using Artificial Neural Networks. In *Dynamic Data Assimilation - Beating the Uncertainties*. IntechOpen. https://doi.org/10.5772/INTECHOPEN.91935

Alpaydin, E. (2014). Design and analysis of machine learning experiments. In *Introduction to Machine Learning* (Issue chapter 11, pp. 547–591). MIT Press.

Alpaydin, E. (2020). *Introduction to machine learning*. The MIT Press. https://mitpress.mit.edu/books/introduction-machine-learning-fourth-edition

Alsop, B., & Elliffe, D. (1988). Concurrent-schedule performance: Effects of relative and overall reinforcer rate. *Journal of the Experimental Analysis of Behavior*, *49*(1), 21–36. https://doi.org/10.1901/jeab.1988.49-21

Alzubaidi, L., Zhang, J., Humaidi, A. J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., Santamaría, J., Fadhel, M. A., Al-Amidie, M., & Farhan, L. (2021). Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. *Journal of Big Data 2021 8:1*, *8*(1), 1–74. https://doi.org/10.1186/S40537-021-00444-8

Ardalan, Z., & Subbian, V. (2022). Transfer learning approaches for neuroimaging analysis: A scoping review. *Frontiers in Artificial Intelligence*, *5*(780405). https://doi.org/10.3389/frai.2022.780405

Aroyehun, S. T., Angel, J., Majumder, N., Gelbukh, A., & Hussain, A. (2021). Leveraging label hierarchy using transfer and multi-task learning: A case study on patent classification. *Neurocomputing*, *464*, 421–431. https://doi.org/10.1016/J.NEUCOM.2021.07.057

Artificial intelligence - Reasoning | Britannica. (n.d.). Retrieved April 26, 2023, from https://www.britannica.com/technology/artificial-intelligence/Reasoning

Asif, M., Martiniano, H. F. M. C., Marques, A. R., Santos, J. X., Vilela, J., Rasga, C., Oliveira, G., Couto, F. M., & Vicente, A. M. (2020). Identification of biological mechanisms underlying a multidimensional ASD phenotype using machine learning. *Translational Psychiatry*, *10*(1), 1–12. https://doi.org/10.1038/s41398-020-0721-1

Auge, D., Hille, J., Mueller, E., & Knoll, A. (2021). A Survey of Encoding Techniques for Signal Processing in Spiking Neural Networks. *Neural Processing Letters*, *53*(6), 4693–4710. https://doi.org/10.1007/S11063-021-10562-2/TABLES/1

BACB - Behavior Analyst Certification Board. (n.d.). Retrieved May 26, 2023, from https://www.bacb.com/

Badillo, S., Banfai, B., Birzele, F., Davydov, I. I., Hutchinson, L., Kam-Thong, T., Siebourg-Polster, J., Steiert, B., & Zhang, J. D. (2020). An introduction to machine learning. *Clinical Pharmacology and Therapeutics*, *107*(4), 871–885. https://doi.org/10.1002/cpt.1796

Baer, D. M. (1976). The organism as host. *Human Development*, *19*(2), 87–98. https://doi.org/10.1159/000271519

Baer, D. M., Wolf, M. M., & Risley, T. R. (1968). Some current dimensions of applied behavior analysis. *Journal of Applied Behavior Analysis*, *1*(1), 91–97. https://doi.org/10.1901/jaba.1968.1-91

Bailey, J. D., Baker, C., Rzeszutek, M. J., & Lanovaz, M. J. (2021). *Machine learning for supplementing behavioral assessment*. 605–619.

Barbiero, P., Squillero, G., & Tonda, A. (2020). *Modeling generalization in machine learning: A methodological and computational study*. *September 2018*, 1–21. http://arxiv.org/abs/2006.15680

Baum, W. M. (1973). The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior*, *20*(1), 137. https://doi.org/10.1901/JEAB.1973.20-137

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, *22*(1), 231–242. https://doi.org/10.1901/jeab.1974.22-231

Baum, W. M. (2002). From molecular to molar: A paradigm shift in behavior analysis. *Journal of the Experimental Analysis of Behavior*, *78*(1), 95–116. https://doi.org/10.1901/jeab.2002.78-95

Baum, W. M. (2010). Dynamics of choice: A tutorial. *Journal of the Experimental Analysis of Behavior*, *94*(2), 161–174. https://doi.org/10.1901/jeab.2010.94-161

Baum, W. M. (2012). Rethinking reinforcement: Allocation, induction, and contingency. *Journal of the Experimental Analysis of Behavior*, *97*(1), 101–124. https://doi.org/10.1901/jeab.2012.97-101

Baum, W. M. (2018). Three laws of behavior: Allocation, induction, and covariance. *Behavior Analysis: Research and Practice*, *18*(3), 239. https://doi.org/10.1037/BAR0000104

Baum, W. M., & William, B. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior*, *32*(2), 269–281. https://doi.org/10.1901/jeab.1979.32-269

Belke, T. W., & Heyman, G. M. (1994). A matching law analysis of the reinforcing efficacy of wheel running in rats. *Animal Learning & Behavior*, *22*(3), 267–274. https://doi.org/10.3758/BF03209834

Bengio, Y., Simard, P., & Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, *5*(2), 157–166. https://doi.org/10.1109/72.279181

Beniaguev, D., Segev, I., & London, M. (2021). Single cortical neurons as deep artificial neural networks. *Neuron*, *109*(17), 2727-2739.e3. https://doi.org/10.1016/J.NEURON.2021.07.002

Bhangale, K. B., & Mohanaprasad, K. (2021). A review on speech processing using machine learning paradigm. *International Journal of Speech Technology 2021 24:2*, *24*(2), 367–388. https://doi.org/10.1007/S10772-021-09808-0

Boné, R., & Cardot, H. (2005). Time delay learning by gradient descent in Recurrent Neural Networks. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *3697 LNCS*, 175–180. https://doi.org/10.1007/11550907_29/COVER

Bouton, M. E., & Balleine, B. W. (2019). Prediction and control of operant behavior: What you see is not all there is. *Behavior Analysis: Research and Practice*, *19*(2), 202–212. https://doi.org/10.1037/bar0000108

Boutros, N., Elliffe, D., & Davison, M. (2010). Time versus response indices affect conclusions about preference pulses. *Behavioural Processes*, *84*(1), 450–454. https://doi.org/10.1016/j.beproc.2009.11.007

Breen, V., Kasabov, N., Du, P., & Calder, S. (2016). A spiking neural network for personalised modelling of electrogastrography (EGG). *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *9896 LNAI*, 18–25. https://doi.org/10.1007/978-3-319-46182-3_2/FIGURES/4

Breland, K., & Breland, M. (1961). The misbehavior of organisms. *American Psychologist*, *16*(11), 681–684. https://doi.org/10.1037/H0040090

Brodhead, M. T., & Truckenmiller, A. (2021). Proof-of-concept evaluation of a decision-making algorithm. *Behavior Analysis: Research and Practice*, *21*(1), 84–89.

Burgos, J. E. (2003). Theoretical note: Simulating latent inhibition with selection neural networks. *Behavioural Processes*, *62*(1–3), 183–192. https://doi.org/10.1016/S0376-6357(03)00025-1

Burgos, J. E. (2007). Autoshaping and automaintenance: A neural-network approach. *Journal of the Experimental Analysis of Behavior*, *88*(1), 115–130. https://doi.org/10.1901/jeab.2007.75-04

Burgos, J. E. (2010). The operant/respondent distinction: A computational neural-network analysis. *Computational Models of Conditioning*, 244–271. https://doi.org/10.1017/CBO9780511760402.009

Burgos, J. E. (2015). Misbehavior in a neural network model. *International Journal of Comparative Psychology*, 28. https://doi.org/10.46867/ijcp.2015.28.01.02

Burgos, J. E. (2018). Is a nervous system necessary for learning? *Perspectives on Behavior Science*, *41*(2), 343–368. https://doi.org/10.1007/s40614-018-00179-7

Burkitt, A. N. (2006). A review of the integrate-and-fire neuron model: I. Homogeneous synaptic input. *Biological Cybernetics*, *95*(1), 1–19. https://doi.org/10.1007/S00422-006-0068-6/METRICS

Burney, V., Arnold-Saritepe, A., & McCann, C. M. (2023). Rethinking the place of qualitative methods in behavior analysis. *Perspectives on Behavior Science*, *46*(1), 185–200. https://doi.org/10.1007/s40614-022-00362-x

Cai, L., Gao, J., & Zhao, D. (2020). A review of the application of deep learning in medical image classification and segmentation. *Annals of Translational Medicine*, *8*(11), 713–713. https://doi.org/10.21037/ATM.2020.02.44

Cantin-Garside, K. D., Kong, Z., White, S. W., Antezana, L., Kim, S., & Nussbaum, M. A. (2020). Detecting and classifying self-injurious behavior in autism spectrum disorder using machine learning techniques. *Journal of Autism and Developmental Disorders*, *50*(11), 4039–4052. https://doi.org/10.1007/s10803-020-04463-x

Cavus, N., Lawan, A. A., Ibrahim, Z., Dahiru, A., Tahir, S., Abdulrazak, U. I., & Hussaini, A. (2021). A systematic literature review on the application of machine-learning models in behavioral assessment of autism spectrum disorder. *Journal of Personalized Medicine*, *11*(4), 299. https://doi.org/10.3390/JPM11040299

Cawley, G. C., & Talbot, N. L. C. (2010). On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, *11*, 2079–2107.

Cearns, M., Hahn, T., & Baune, B. T. (2019). Recommendations and future directions for supervised machine learning in psychiatry. *Translational Psychiatry 2019 9:1*, *9*(1), 1–12. https://doi.org/10.1038/s41398-019-0607-2

Chen, X., Kang, W., Zhu, D., Zhang, X., Lei, N., Zhang, Y., Zhou, Y., & Zhao, W. (2018). A compact skyrmionic leaky-integrate-fire spiking neuron device. *Nanoscale*, *10*(13), 6139–6146. https://doi.org/10.1039/c7nr09722k

Chiesa, M. (1994). *Radical behaviorism: The philosophy and the science*. Boston: Authors Cooperative.

Chung, J., Gulcehre, C., Cho, K., & Bengio, Y. (2014). *Empirical evaluation of gated recurrent neural networks on sequence modeling*. https://doi.org/10.48550/arxiv.1412.3555

Collin, C. B., Gebhardt, T., Golebiewski, M., Karaderi, T., Hillemanns, M., Khan, F. M., Salehzadeh-Yazdi, A., Kirschner, M., Krobitsch, S., & Kuepfer, L. (2022). Computational models for clinical applications in personalized medicine-guidelines and recommendations for Data Integration and Model Validation. *Journal of Personalized Medicine 2022, Vol. 12, Page 166*, *12*(2), 166. https://doi.org/10.3390/JPM12020166

Cortes, C., Vapnik, V., & Saitta, L. (1995). Support-vector networks. *Machine Learning 1995 20:3*, *20*(3), 273–297. https://doi.org/10.1007/BF00994018

Cowie, S. (2018). Behavioral time travel: Control by past, present, and potential events. *Behavior Analysis: Research and Practice*, *18*(2), 174. https://doi.org/10.1037/BAR0000122

Cowie, S. (2019). Some weaknesses of a response-strength account of reinforcer effects. *European Journal of Behavior Analysis*, *22*(2), 1–16. https://doi.org/10.1080/15021149.2019.1685247

Cowie, S., & Davison, M. (2016a). Control by reinforcers across time and space: A review of recent choice research. *Journal of the Experimental Analysis of Behavior*, *105*(2), 246–269. https://doi.org/10.1002/JEAB.200

Cowie, S., & Davison, M. (2016b). Control by reinforcers across time and space: A review of recent choice research. *Journal of the Experimental Analysis of Behavior*, *105*(2), 246–269. https://doi.org/10.1002/jeab.200

Cowie, S., & Davison, M. (2020a). Being there on time: Reinforcer effects on timing and locating. *Journal of the Experimental Analysis of Behavior*, *113*(2), 340–362. https://doi.org/10.1002/JEAB.581

Cowie, S., & Davison, M. (2020b). Generalizing from the past, choosing the future. *Perspectives on Behavior Science*, *43*(2), 245. https://doi.org/10.1007/S40614-020-00257-9

Cowie, S., Davison, M., & Elliffe, D. (2011). Reinforcement: Food signals the time and location of future food. *Journal of the Experimental Analysis of Behavior*, *96*(1), 63–86. https://doi.org/10.1901/jeab.2011.96-63

Cowie, S., Davison, M., & Elliffe, D. (2014). A model for food and stimulus changes that signal time-based contingency changes. *Journal of the Experimental Analysis of Behavior*, *102*(3), 289–310. https://doi.org/10.1002/jeab.105

Cowie, S., Davison, M., & Elliffe, D. (2016). A model for discriminating reinforcers in time and space. *Behavioural Processes*, *127*, 62–73. https://doi.org/10.1016/j.beproc.2016.03.010

Cowie, S., Elliffe, D., & Davison, M. (2013). Concurrent schedules: Discriminating reinforcer-ratio reversals at a fixed time after the previous reinforcer. *Journal of the Experimental Analysis of Behavior*, *100*(2), 117–134. https://doi.org/10.1002/jeab.43

Cowie, S., Virués-Ortega, J., McCormack, J., Hogg, P., & Podlesnik, C. A. (2021). Extending a misallocation model to children’s choice behavior. *Journal of Experimental Psychology. Animal Learning and Cognition*, *47*(3), 317–325. https://doi.org/10.1037/XAN0000299

Cox, D. J. (2019). The many functions of quantitative modeling. *Computational Brain & Behavior*, *2*(3–4), 166–169. https://doi.org/10.1007/s42113-019-00048-9

Cox, D. J., Klapes, B., & Falligant, J. M. (2021). Scaling N from 1 to 1,000,000: Application of the generalized matching law to big data contexts. *Perspectives on Behavior Science*, *44*(4), 641–665. https://doi.org/10.1007/S40614-021-00298-8

Critchfield, T. S. (2011). Translational contributions of the experimental analysis of behavior. *The Behavior Analyst*, *34*(1), 3. https://doi.org/10.1007/BF03392227

Cyr, A., Boukadoum, M., & Thériault, F. (2014). Operant conditioning: A minimal components requirement in artificial spiking neurons designed for bio-inspired Robot's controller. *Frontiers in Neurorobotics*, *8*(JUL), 1–32. https://doi.org/10.3389/fnbot.2014.00021

Cyr, A., Morand-Ferron, J., & Thériault, F. (2020). Dual exploration strategies using artificial spiking neural networks in a robotic learning task. *International Journal for Adaptive Behavior*, *29*(6), 567–578. https://doi.org/10.1177/1059712320924744

Davidson, S., & Furber, S. B. (2021). Comparison of artificial and spiking neural networks on digital hardware. *Frontiers in Neuroscience*, *15*(April), 1–7. https://doi.org/10.3389/fnins.2021.651141

Davison, M., & Baum, W. M. (2000). Choice in a variable environment: Every reinforcer counts. *Journal of the Experimental Analysis of Behavior*, *74*(1), 1–24. https://doi.org/10.1901/jeab.2000.74-1

Davison, M., & Baum, W. M. (2002). Choice in a variable environment: Effects of blackout duration and extinction between components. *Journal of the Experimental Analysis of Behavior*, *77*(1), 65–89. https://doi.org/10.1901/JEAB.2002.77-65

Davison, M., & Baum, W. M. (2003). Every reinforcer counts: Reinforcer magnitude and local preference. *Journal of the Experimental Analysis of Behavior*, *80*(1), 95–129. https://doi.org/10.1901/JEAB.2003.80-95

Davison, M., & Cowie, S. (2022). Modeling choice across time: Effects of response–reinforcer discriminability. *Journal of the Experimental Analysis of Behavior*, *117*(1), 36–52. https://doi.org/10.1002/JEAB.723

Davison, M., & Elliffe, D. (2009). Variance matters: The shape of a datum. *Behavioural Processes*, *81*(2), 216–222. https://doi.org/10.1016/j.beproc.2009.01.004

Davison, M., & Jenkins, P. E. (1985). Stimulus discriminability, contingency discriminability, and schedule performance. *Animal Learning & Behavior*, *13*(1), 77–84. https://doi.org/10.3758/BF03213368

Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Lawrence Erlbaum Associates Inc. https://doi.org/https://doi.org/10.4324/9781315638911

Dehghan, P., Alashwal, H., & Moustafa, A. A. (2022). Applications of machine learning to behavioral sciences: focus on categorical data. *Discover Psychology*, *2*(1). https://doi.org/10.1007/s44202-022-00027-5

Dezfouli, A., Griffiths, K., Ramos, F., Dayan, P., & Balleine, B. W. (2019). Models that learn how humans learn: The case of decision-making and its disorders. In *PLoS Computational Biology* (Vol. 15, Issue 6). https://doi.org/10.1371/journal.pcbi.1006903

DiPietro, R., & Hager, G. D. (2020). Deep learning: RNNs and LSTM. *Handbook of Medical Image Computing and Computer Assisted Intervention*, 503–519. https://doi.org/10.1016/B978-0-12-816176-0.00026-0

Doborjeh, M., Doborjeh, Z., Merkin, A., Krishnamurthi, R., Enayatollahi, R., Feigin, V., & Kasabov, N. (2022). Personalized spiking neural network models of clinical and environmental factors to predict stroke. *Cognitive Computation*, *14*(6), 2187–2202. https://doi.org/10.1007/S12559-021-09975-X/FIGURES/9

Doborjeh, M., Kasabov, N., Doborjeh, Z., Enayatollahi, R., Tu, E., & Gandomi, A. H. (2019). Personalised modelling with spiking neural networks integrating temporal and static information. *Neural Networks*, *119*, 162–177. https://doi.org/10.1016/J.NEUNET.2019.07.021

Dudani, S. A. (1976). The Distance-Weighted k-Nearest-Neighbor Rule. *IEEE Transactions on Systems, Man and Cybernetics*, *SMC-6*(4), 325–327. https://doi.org/10.1109/TSMC.1976.5408784

Dufour, M. M., Lanovaz, M. J., & Cardinal, P. (2020). Artificial intelligence for the measurement of vocal stereotypy. *Journal of the Experimental Analysis of Behavior*, *114*(3), 368–380. https://doi.org/10.1002/jeab.636

Dwyer, D. B., Falkai, P., & Koutsouleris, N. (2018). Machine learning approaches for clinical psychology and psychiatry. *Annual Review of Clinical Psychology*, *14*(1), 91–118. https://doi.org/10.1146/annurev-clinpsy-032816-045037

Ebbehoj, A., Thunbo, M. Ø., Andersen, O. E., Glindtvad, M. V., & Hulman, A. (2022). Transfer learning for non-image data in clinical research: A scoping review. *PLOS Digital Health*, *1*(2), e0000014. https://doi.org/10.1371/journal.pdig.0000014

Ehrmann, D. E., Joshi, S., Goodfellow, S. D., Mazwi, M. L., & Eytan, D. (2023). Making machine learning matter to clinicians: model actionability in medical decision-making. *Npj Digital Medicine 2023 6:1*, *6*(1), 1–5. https://doi.org/10.1038/s41746-023-00753-7

Fan, H., & Poole, M. S. (2006). What Is personalization? Perspectives on the design and implementation of personalization in information systems. *Journal of*

Organizational Computing and Electronic Commerce, *16*(3–4), 179–202. https://doi.org/10.1080/10919392.2006.9681199

Fan, X., & Markram, H. (2019). A brief history of simulation neuroscience. *Frontiers in Neuroinformatics*, *13*(May), 1–28. https://doi.org/10.3389/fninf.2019.00032

Freeman, T. J., & Lattal, K. A. (1992). Stimulus control of behavioral history. *Journal of the Experimental Analysis of Behavior*, *57*(1), 5–15. https://doi.org/10.1901/JEAB.1992.57-5

Fu, Q., & Dong, H. (2021). An ensemble unsupervised spiking neural network for objective recognition. *Neurocomputing*, *419*, 47–58. https://doi.org/10.1016/j.neucom.2020.07.109

Georgioudakis, M., & Plevris, V. (2020). A comparative study of differential evolution variants in constrained structural optimization. *Frontiers in Built Environment*, *6*, 102. https://doi.org/10.3389/FBUIL.2020.00102/BIBTEX

Gerstner, W., & Kistler, W. M. (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge University Press. https://doi.org/doi.org/10.1017/CBO9780511815706

Gerstner, W., Kistler, W. M., Naud, R., & Paninski, L. (2014). Neuronal dynamics: From single neurons to networks and models of cognition. In *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press. https://doi.org/10.1017/CBO9781107447615

Ghosh-Dastidar, S., & Adeli, H. (2009). Spiking neural networks. *International Journal of Neural Systems*, *19*(4), 295–308. https://doi.org/10.1142/S0129065709002002

Gomes-Ng, S., Elliffe, D., & Cowie, S. (2017). How do reinforcers affect choice? Preference pulses after responses and reinforcers. *Journal of the Experimental Analysis of Behavior*, *108*(1), 17–38. https://doi.org/10.1002/JEAB.260

Graves, A. (2012). *Supervised sequence labeling with recurrent neural networks* (Vol. 385). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-24797-2

Han, J., Kamber, M., & Pei, J. (2012). Cluster analysis: Basic concepts and methods. In J. Han, M. Kamber, & J. Pei (Eds.), *Data Mining (Third Edition)* (pp. 443–495). Morgan Kaufmann. https://doi.org/https://doi.org/10.1016/B978-0-12-381479-1.00010-1

Hayes, S. C. (1981). Single case experimental design and empirical clinical practice. *Journal of Consulting and Clinical Psychology*, *49*(2), 193–211. https://doi.org/10.1037/0022-006X.49.2.193

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, *4*(3), 267–272. https://doi.org/10.1901/jeab.1961.4-267

Herrnstein, R. J. (1970). On the law of effect. *Journal of Experimental Analysis of Behavior*, *13*(2), 243–246. https://doi.org/10.1901/jeab.1970.13-243

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Hollard, V., & Davison, M. C. (1971). Preference for qualitatively different reinforcers. *Journal of the Experimental Analysis of Behavior*, *16*(3), 375–380. https://doi.org/10.1901/jeab.1971.16-375

Hoshino, T., Kanoga, S., Tsubaki, M., & Aoyama, A. (2022). Comparing subject-to-subject transfer learning methods in surface electromyogram-based motion recognition with shallow and deep classifiers. *Neurocomputing*. https://doi.org/10.1016/J.NEUCOM.2021.12.081

Hunter, I., & Davison, M. (1985). Determination of a behavioral transfer function: White-noise analysis of session-to-session response-ratio dynamics on concurrent

VI VI schedules. *Journal of the Experimental Analysis of Behavior*, *43*(1), 43. https://doi.org/10.1901/JEAB.1985.43-43

Imam, A. A. (2018). Place of behavior analysis in the changing culture of replication and statistical reporting in psychological science. *European Journal of Behavior Analysis*, *19*(1), 2–10. https://doi.org/10.1080/15021149.2018.1463123

Imam, A. A., & Frate, M. (2019). A snapshot look at replication and statistical reporting practices in psychology journals. *European Journal of Behavior Analysis*, *20*(2), 204–229. https://doi.org/10.1080/15021149.2019.1680179

Iman, M., Rasheed, K., & Arabnia, H. R. (2023). A review of deep transfer learning and recent advancements. *Technologies*, *11*(40). https://doi.org/10.3390/technologies11020040

Ioannidis, J. P. A. (2016). Why most clinical research is not useful. *PLoS Medicine*, *13*(6), 1–10. https://doi.org/10.1371/journal.pmed.1002049

Ioannidis, J. P. A. (2018). Why most published research findings are false. *Getting to Good: Research Integrity in the Biomedical Sciences*, *2*(8), 2–8. https://doi.org/10.1371/journal.pmed.0020124

Iwata, B. A., Dorsey, M. F., Slifer, K. J., Bauman, K. E., & Richman, G. S. (1994). Towards a functional analysis of self-injury. *Journal of Applied Behavior Analysis*, *27*(2), 197–209. https://doi.org/10.1901/jaba.1994.27-197

Izhikevich, E. M. (2004). Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks*, *15*(5), 1063–1070. https://doi.org/10.1109/TNN.2004.832719

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning: with Applications in R*. Springer. https://doi.org/10.1016/j.peva.2007.06.006

Johnson, J. M., & Khoshgoftaar, T. M. (2019). Survey on deep learning with class imbalance. *Journal of Big Data*, *6*(1), 1–54. https://doi.org/10.1186/S40537-019-0192-5/TABLES/18

Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, *349*(6245), 255–260. https://doi.org/DOI: 10.1126/science.aaa84

Justice, L. M., Ahn, W. Y., & Logan, J. A. R. (2019). Identifying children with clinical language disorder: An application of machine-learning classification. *Journal of Learning Disabilities*, *52*(5), 351–365. https://doi.org/10.1177/0022219419845070

Karekla, M., Demosthenous, G., Georgiou, C., Konstantinou, P., Trigiorgi, A., Koushiou, M., Panayiotou, G., & Gloster, A. T. (2022). Machine learning advances the classification and prediction of responding from psychophysiological reactions. *Journal of Contextual Behavioral Science*, *26*(November 2021), 36–43. https://doi.org/10.1016/j.jcbs.2022.07.006

Kasabov, N. (2019). Evolving and spiking connectionist systems for brain-inspired artificial intelligence. *Artificial Intelligence in the Age of Neural Networks and Brain Computing*, 111–138. https://doi.org/10.1016/B978-0-12-815480-9.00006-2

Kasabov, N. (2014). NeuCube: a spiking neural network architecture for mapping, learning and understanding of spatio-temporal brain data. *Neural Networks: The Official Journal of the International Neural Network Society*, *52*, 62–76. https://doi.org/10.1016/J.NEUNET.2014.01.006

Kasabov, N. K. (2019). *A computational framework for personalised modelling. Applications in bioinformatics*. 563–591. https://doi.org/10.1007/978-3-662-57715-8_17

Kassraian-Fard, P., Matthis, C., Balsters, J. H., Maathuis, M. H., & Wenderoth, N. (2016). Promises, pitfalls, and basic guidelines for applying machine learning classifiers to psychiatric imaging data, with autism as an example. *Frontiers in Psychiatry*, *7*(DEC), 177. https://doi.org/10.3389/FPSYT.2016.00177/BIBTEX

Kazdin, A. E. (2011). *Single-case research designs: methods for clinical and applied settings* (2nd ed..). New York: Oxford University Press.

Kelley, M. E., Nadler, C. B., Rey, C., Cowie, S., & Podlesnik, C. A. (2017). Noncontingent reinforcement competes with response performance. *Journal of the Experimental Analysis of Behavior*, *107*(3), 343–353. https://doi.org/10.1002/JEAB.255

Khadeer, A. (2020). Brain-inspired spiking neural networks. *Biomimetics*. https://doi.org/10.5772/INTECHOPEN.93435

Khurana, D., Koli, A., Khatter, K., & Singh, S. (2023). Natural language processing: State of the art, current trends and challenges. *Multimedia Tools and Applications*, *82*(3), 3713–3744. https://doi.org/10.1007/s11042-022-13428-4

Killeen, P. (1972). The matching law. *Journal of the Experimental Analysis of Behavior*, *17*(3), 489–495. https://doi.org/10.1901/jeab.1972.17-489

Killeen, P. R. (1978). Stability criteria. *Journal of the Experimental Analysis of Behavior*, *29*(1), 17. https://doi.org/10.1901/JEAB.1978.29-17

Killeen, P. R. (2015). The logistics of choice. *Journal of the Experimental Analysis of Behavior*, *104*(1), 74–92. https://doi.org/10.1002/jeab.156

Killeen, P. R. (2018). The futures of experimental analysis of behavior. *Behavior Analysis: Research and Practice*, *18*(2), 124–133. https://doi.org/10.1037/BAR0000100

Killeen, P. R. (2019). Predict, control, and replicate to understand: How statistics can foster the fundamental goals of science. *Perspectives on Behavior Science*, *42*(1), 109–132. https://doi.org/10.1007/s40614-018-0171-8

Killeen, P. R., & Jacobs, K. W. (2016). Coal is not black, snow is not white, food is not a reinforcer: The roles of affordances and dispositions in the analysis of behavior. *The Behavior Analyst*, *40*(1), 17–38. https://doi.org/10.1007/S40614-016-0080-7

Kiranyaz, S., Avci, O., Abdeljaber, O., Ince, T., Gabbouj, M., & Inman, D. J. (2021). 1D convolutional neural networks and applications: A survey. *Mechanical Systems and Signal Processing*, *151*, 107398. https://doi.org/10.1016/J.YMSSP.2020.107398

Koppe, G., Meyer-Lindenberg, A., & Durstewitz, D. (2020). Deep learning for small and big data in psychiatry. *Neuropsychopharmacology 2020 46:1*, *46*(1), 176–190. https://doi.org/10.1038/s41386-020-0767-z

Koppe, G., Meyer-Lindenberg, A., & Durstewitz, D. (2021). Deep learning for small and big data in psychiatry. *Neuropsychopharmacology*, *46*(1), 176–190. https://doi.org/10.1038/s41386-020-0767-z

Korkmaz, C., & Correia, A. P. (2019a). A review of research on machine learning in educational technology. *Educational Media International*, *56*(3), 250–267. https://doi.org/10.1080/09523987.2019.1669875

Korkmaz, C., & Correia, A. P. (2019b). A review of research on machine learning in educational technology. *Educational Media International*, *56*(3), 250–267. https://doi.org/10.1080/09523987.2019.1669875

Krägeloh, C. U., Davison, M., & Elliffe, D. M. (2005). Local preference in concurrent schedules: The effects of reinforcer sequences. *Journal of the Experimental Analysis of Behavior*, *84*(1), 37–64. https://doi.org/10.1901/jeab.2005.114-04

Krägeloh, C. U., Zapanta, A. E., Shepherd, D., & Landon, J. (2010). Human choice behaviour in a frequently changing environment. *Behavioural Processes*, *83*(1), 119–126. https://doi.org/10.1016/j.beproc.2009.11.005

Krebs, C., Weinberg, J., Akesson, E., & Dilli, E. (2018). *Neuroscience* (2nd ed.). Wolters Kluwer Health.

Kuhn, M., & Johnson, K. (2018). *Applied predictive modeling*. Springer. https://doi.org/10.1007/978-1-4614-6849-3

Landon, J., & Davison, M. (2001). Reinforcer-ratio variation and its effects on rate of adaptation. *Journal of the Experimental Analysis of Behavior*, *75*(2), 207–234. https://doi.org/10.1901/jeab.2001.75-207

Landon, J., Davison, M., & Elliffe, D. (2002). Concurrent schedules: Short-and long-term effects of reinforcers. *Journal of the Experimental Analysis of Behavior*, *77*(3), 257–271. https://doi.org/10.1901/JEAB.2002.77-257

Landon, J., Davison, M., & Elliffe, D. (2003a). Choice in a variable environment: Effects of unequal reinforcer distributions. *Journal of the Experimental Analysis of Behavior*, *80*(2), 187–204. https://doi.org/10.1901/JEAB.2003.80-187

Landon, J., Davison, M., & Elliffe, D. (2003b). Concurrent schedules: Reinforcer magnitude effects. *Journal of the Experimental Analysis of Behavior*, *79*(3), 351–365. https://doi.org/10.1901/jeab.2003.79-351

Landon, J., Davison, M., Krägeloh, C. U., Thompson, N. M., Miles, J. L., Vickers, M. H., Fraser, M., & Breier, B. H. (2007). Global undernutrition during gestation. *Learning & Behavior*, *35*(2), 79–86.

Lanovaz, M. J., Giannakakos, A. R., & Destras, O. (2020). Machine learning to analyze single-case data: A proof of concept. *Perspectives on Behavior Science*, *43*(1), 21–38. https://doi.org/10.1007/s40614-020-00244-0

Lattal, K. A. (1995). Contingency and behavior analysis. *The Behavior Analyst*, *18*(2), 209. https://doi.org/10.1007/BF03392709

LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998a). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, *86*(11), 2278–2323. https://doi.org/10.1109/5.726791

LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998b). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, *86*(11), 2278–2323. https://doi.org/10.1109/5.726791

Li, Z., Liu, F., Yang, W., Peng, S., & Zhou, J. (2022). A survey of convolutional neural networks: Analysis, applications, and prospects. *IEEE Transactions on Neural Networks and Learning Systems*, *33*(12), 6999–7019. https://doi.org/10.1109/TNNLS.2021.3084827

Lin, J.-W., Chang, Y., Chang, C.-H., & Ko, L.-W. (2021). Application of transfer learning in field verification for children in Attention Deficit Hyperactivity Disorder. *2021 International Conference on Fuzzy Theory and Its Applications (IFUZZY)*, 1–4. https://doi.org/10.1109/iFUZZY53132.2021.9605078

Linstead, E., Dixon, D. R., French, R., Granpeesheh, D., Adams, H., German, R., Powell, A., Stevens, E., Tarbox, J., & Kornack, J. (2017). Intensity and learning outcomes in the treatment of children with Autism Spectrum Disorder. *Behavior Modification*, *41*(2), 229–252. https://doi.org/10.1177/0145445516667059

Linstead, E., German, R., Dixon, D., Granpeesheh, D., Novack, M., & Powell, A. (2015). An application of neural networks to predicting mastery of learning outcomes in the treatment of Autism Spectrum Disorder. *2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA)*, 414–418. https://doi.org/10.1109/ICMLA.2015.214

Lobb, B., & Davison, M. C. (1975). Performance in concurrent interval schedules: a systematic replication. *Journal of the Experimental Analysis of Behavior*, *24*(2), 191–197. https://doi.org/10.1901/JEAB.1975.24-191

Locey, M. L. (2020). The evolution of behavior analysis: Toward a replication crisis? *Perspectives on Behavior Science*, *43*(4), 655–675. https://doi.org/10.1007/s40614-020-00264-w

Lu, X. (2022). Deep learning based emotion recognition and visualization of figural representation. *Frontiers in Psychology*, *12*, 6237. https://doi.org/10.3389/FPSYG.2021.818833/BIBTEX

Maass, W. (1997). Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, *10*(9), 1659–1671. https://doi.org/10.1016/S0893-6080(97)00011-7

Maass, W., & Markram, H. (2004). On the computational power of circuits of spiking neurons. *Journal of Computer and System Sciences*, *69*(4), 593–616. https://doi.org/10.1016/j.jcss.2004.04.001

Macpherson, T., Churchland, A., Sejnowski, T., DiCarlo, J., Kamitani, Y., Takahashi, H., & Hikida, T. (2021). Natural and Artificial Intelligence: A brief introduction to the interplay between AI and neuroscience research. *Neural Networks*. https://doi.org/10.1016/j.neunet.2021.09.018

Maleki, F., Muthukrishnan, N., Ovens, K., Reinhold, C., & Forghani, R. (2020). Machine learning algorithm validation: From essentials to advanced applications and implications for regulatory certification and deployment. *Neuroimaging Clinics of North America*, *30*(4), 433–445. https://doi.org/10.1016/J.NIC.2020.08.004

Malgaroli, M., & Schultebraucks, K. (2020). Artificial intelligence and Posttraumatic Stress Disorder (PTSD). *European Psychologist*, *25*(4), 272–282. https://doi.org/10.1027/1016-9040/a000423

Mazur, J. E. (1997). Effects of rate of reinforcement and rate of change on choice behaviour in transition. *The Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology*, *50*(2), 111–128. https://doi.org/10.1080/713932646

Mazur, J. E. (2006). Mathematical models and the experimental analysis of behavior. *Journal of the Experimental Analysis of Behavior*, *85*(2), 275–291. https://doi.org/10.1901/JEAB.2006.65-05

Mazur, J. E. (2016). Choice. In *Learning and Behavior* (8th ed., pp. 37–51). Routledge. https://doi.org/https://doi.org/10.4324/9781315450285

McLean, A. P., Grace, R. C., Pitts, R. C., & Hughes, C. E. (2014). Preference pulses without reinforcers. *Journal of the Experimental Analysis of Behavior*, *101*(3), 317–336. https://doi.org/10.1002/JEAB.84

Meehan, A. J., Lewis, S. J., Fazel, S., Fusar-Poli, P., Steyerberg, E. W., Stahl, D., & Danese, A. (2022). Clinical prediction models in psychiatry: a systematic review of two decades of progress and challenges. *Molecular Psychiatry*, *27*(6), 2700–2708. https://doi.org/10.1038/S41380-022-01528-4

Miranda-Dukoski, L., Davison, M., & Elliffe, D. (2014). Choice, time and food: Continuous cyclical changes in food probability between reinforcers. *Journal of the Experimental Analysis of Behavior*, *101*(3), 406–421. https://doi.org/10.1002/jeab.79

Mostafa, S., & Wu, F. X. (2021). Diagnosis of autism spectrum disorder with convolutional autoencoder and structural MRI images. *Neural Engineering*

Techniques for Autism Spectrum Disorder: Volume 1: Imaging and Signal Analysis, 23–38. https://doi.org/10.1016/B978-0-12-822822-7.00003-X

Ninness, C., & Ninness, S. K. (2020). Emergent virtual analytics: Modeling contextual control of derived stimulus relations. *Behavior and Social Issues*, *29*(1), 119–137. https://doi.org/10.1007/S42822-020-00032-0

Ninness, C., Ninness, S. K., Rumph, M., & Lawson, D. (2018). The emergence of stimulus relations: Human and computer learning. *Perspectives on Behavior Science*, *41*(1), 121–154. https://doi.org/10.1007/s40614-017-0125-6

Ninness, C., Rehfeldt, R. A., & Ninness, S. K. (2019). Identifying accurate and inaccurate stimulus relations: Human and computer learning. *Psychological Record*, *69*(3), 333–356. https://doi.org/10.1007/s40732-019-00337-6

Niu, S., Liu, Y., Wang, J., & Song, H. (2021). A decade survey of transfer learning (2010–2020). *IEEE Transactions on Artificial Intelligence*, *1*(2), 151–166. https://doi.org/10.1109/tai.2021.3054609

Okouchi, H., & Lattal, K. A. (2006). An analysis of reinforcement history effects. *Journal of the Experimental Analysis of Behavior*, *86*(1), 31–42. https://doi.org/10.1901/jeab.2006.75-05

Okouchi, H., Lattal, K. A., Sonoda, A., & Nakamae, T. (2014). Stimulus control and generalization of remote behavioral history. *Journal of the Experimental Analysis of Behavior*, *101*(2), 275–287. https://doi.org/10.1002/jeab.75

Orrù, G., Monaro, M., Conversano, C., Gemignani, A., & Sartori, G. (2020). Machine learning in psychometrics and psychological research. *Frontiers in Psychology*, *10*(January), 1–10. https://doi.org/10.3389/fpsyg.2019.02970

O’Shea, K., & Nash, R. (2015). An introduction to convolutional neural networks. *International Journal for Research in Applied Science and Engineering Technology*, *10*(12), 943–947. https://doi.org/10.22214/ijraset.2022.47789

Overskeid, G. (2018). Do we need the environment to explain operant behavior? *Frontiers in Psychology*, *9*(MAR), 373. https://doi.org/10.3389/FPSYG.2018.00373/BIBTEX

Pfeiffer, M., & Pfeil, T. (2018). Deep learning with spiking neurons: Opportunities and challenges. *Frontiers in Neuroscience*, *12*(October). https://doi.org/10.3389/fnins.2018.00774

Pierce, W. D., & Cheney, C. D. (2004). *Behavior analysis and learning* (3rd edition). Lawrence Erlbaum Associates Publishers.

Pierce, W. D., & Epling, W. F. (1983). Choice, matching, and human behavior: A review of the literature. *The Behavior Analyst*, *6*(1), 57–76. https://doi.org/10.1007/BF03391874

Pipkin, C. St. P., & Vollmer, T. R. (2009). Applied implications of reinforcement history effects. *Journal of Applied Behavior Analysis*, *42*(1), 83–103. https://doi.org/10.1901/JABA.2009.42-83

Plaud, J. J. (1992). The prediction and control of behavior revisited: A review of the matching law. *Journal of Behavior Therapy and Experimental Psychiatry*, *23*(1), 25–31. https://doi.org/10.1016/0005-7916(92)90021-A

Plessas, A., Espinosa-Ramos, J. I., Parry, D., Cowie, S., & Landon, J. (2022). Machine learning with a snapshot of data: Spiking neural network ‘predicts’ reinforcement histories of pigeons’ choice behavior. *Journal of the Experimental Analysis of Behavior*. https://doi.org/10.1002/JEAB.759

Podlesnik, C. A., Ritchey, C. M., Kuroda, T., & Cowie, S. (2022). A quantitative analysis of the effects of alternative reinforcement rate and magnitude on resurgence. *Behavioural Processes*, *198*, 104641. https://doi.org/10.1016/J.BEPROC.2022.104641

Poling, A. (2010). Looking to the future: Will behavior analysis survive and prosper? *Behavior Analyst*, *33*(1), 7–17. https://doi.org/10.1007/BF03392200

Qian, W., He, Z., Chen, C., & Peng, S. (2022). Partner learning: A comprehensive knowledge transfer for vehicle re-identification. *Neurocomputing*, *480*, 89–98. https://doi.org/10.1016/J.NEUCOM.2022.01.043

Rafi, T. H. (2021). *A brief review on spiking neural network - A biological inspiration*. *April*. https://doi.org/10.20944/preprints202104.0202.v1

Ranaweera, M., & Mahmoud, Q. H. (2021). Virtual to real-world transfer learning: A systematic review. In *Electronics* (Vol. 10, Issue 12). https://doi.org/10.3390/electronics10121491

Ratana, R., Sharifzadeh, H., Krishnan, J., & Pang, S. (2019). A comprehensive review of computational methods for automatic prediction of schizophrenia with insight into indigenous populations. *Frontiers in Psychiatry*, *10*(September), 1–15. https://doi.org/10.3389/fpsyt.2019.00659

Rebala, G., Ravi, A., & Churiwala, S. (2019). Machine learning definition and basics. In G. Rebala, A. Ravi, & S. Churiwala (Eds.), *An Introduction to Machine Learning* (pp. 1–17). Springer International Publishing. https://doi.org/10.1007/978-3-030-15729-6_1

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016a). *Model-agnostic interpretability of machine learning*. https://arxiv.org/abs/1606.05386v1

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016b). "Why should I trust you?" Explaining the predictions of any classifier. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, *13-17-Augu*, 1135–1144. https://doi.org/10.1145/2939672.2939778

Riley, R. D., Ensor, J., Snell, K. I. E., Harrell, F. E., Martin, G. P., Reitsma, J. B., Moons, K. G. M., Collins, G., & Van Smeden, M. (2020). Calculating the sample size required for developing a clinical prediction model. *The BMJ*, *368*(March), 1–12. https://doi.org/10.1136/bmj.m441

Roy, K., Jaiswal, A., & Panda, P. (2019). Towards spike-based machine intelligence with neuromorphic computing. *Nature 2019 575:7784*, *575*(7784), 607–617. https://doi.org/10.1038/s41586-019-1677-2

Saeedinia, S. A., Jahed-Motlagh, M. R., Tafakhori, A., & Kasabov, N. (2021). Design of MRI structured spiking neural networks and learning algorithms for personalized modelling, analysis, and prediction of EEG signals. *Scientific Reports 2021 11:1*, *11*(1), 1–14. https://doi.org/10.1038/s41598-021-90029-5

Sarbin, T. R. (1944). The logic of prediction in psychology. *Psychological Review*, *51*(4), 210–228. https://doi.org/10.1037/h0057400

Sarker, I. H. (2021). Machine learning: Algorithms, real-world applications and research directions. *SN Computer Science*, *2*(3), 1–21. https://doi.org/10.1007/S42979-021-00592-X/FIGURES/11

Schlinger, H. D. (2015). Behavior analysis and behavioral neuroscience. *Frontiers in Human Neuroscience*, *9*(APR), 1–4. https://doi.org/10.3389/FNHUM.2015.00210/BIBTEX

Schlinger, H. D. (2018). The heterodoxy of behavior analysis. *Archives of Scientific Psychology*, *6*(1), 159–168. https://doi.org/10.1037/arc0000051

Schneider, J., & Vlachos, M. (2021). Personalization of deep learning. *Data Science – Analytics and Applications*, 89–96. https://doi.org/10.1007/978-3-658-32182-6_14

Schofield, G., & Davison, M. (1997). Nonstable concurrent choice in pigeons. *Journal of the Experimental Analysis of Behavior*, *68*(2), 219–232. https://doi.org/10.1901/JEAB.1997.68-219

Shahan, T. A. (2017). Moving beyond reinforcement and response strength. *Behavior Analyst*, *40*(1), 107–121. https://doi.org/10.1007/s40614-017-0092-y

Shatte, A. B. R., Hutchinson, D. M., & Teague, S. J. (2019). Machine learning in mental health: A scoping review of methods and applications. *Psychological Medicine*. https://doi.org/10.1017/S0033291719000151

Shull, R. L., & Pliskoff, S. S. (1967). Changeover delay and concurrent schedules: some effects on relative performance measures. *Journal of the Experimental Analysis of Behavior*, *10*(6), 517–527. https://doi.org/10.1901/JEAB.1967.10-517

Sidman, Murray. (1960). *Tactics of scientific research*. Basic Books, Inc., Pub.

Simon, C., Bernardy, J. L., & Cowie, S. (2020). On the “strength” of behavior. *Perspectives on Behavior Science*, *43*(4), 677–696. https://doi.org/10.1007/S40614-020-00269-5/FIGURES/2

Skinner, B. F. (1938). *The behavior of organisms*. Appleton-Century-Crofts.

Skinner, B. F. (1966a). The phylogeny and ontogeny of behavior. Contingencies of reinforcement throw light on contingencies of survival in the evolution of behavior. *Science (New York, N.Y.)*, *153*(3741), 1205–1213. https://doi.org/10.1126/science.153.3741.1205

Skinner, B. F. (1966b). What is the experimental analysis of behavior? *Journal of the Experimental Analysis of Behavior*, *9*(3), 213. https://doi.org/10.1901/JEAB.1966.9-213

Skinner, B. F. (1969). *Contingencies of reinforcement a theoretical analysis.* Appleton-Century-Crofts.

Skinner, B. F. (1984a). Methods and theories in the experimental analysis of behavior. *Behavioral and Brain Sciences*, *7*(4), 511–523. https://doi.org/10.1017/S0140525X00026996

Skinner, B. F. (1984b). The operational analysis of psychological terms. *Behavioral and Brain Sciences*, *7*(4), 547–553. https://doi.org/10.1017/S0140525X00027187

Skinner, B. F. (1989). The origins of cognitive thought. *American Psychologist*, *44*(1), 13–18. https://doi.org/10.1037/0003-066X.44.1.13

Skinner, B. F. (Burrhus F. (1981). Selection by consequences. *Science*, *213*(4507), 501–504. https://doi.org/0036-8075/81/0731-0501$01.0

Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. Appleton-Century-Crofts. https://doi.org/https://doi.org/10.1037/10627-000

Soni, D. (2018). *Spiking Neural Networks, the Next Generation of Machine Learning*. https://towardsdatascience.com/spiking-neural-networks-the-next-generation-of-machine-learning-84e167f4eb2b

Sosine, J., & Cox, D. J. (2023). Identifying trends in the open-access behavior analytic literature via computational analyses (I): Simple descriptions of text. *The Analysis of Verbal Behavior*, *I*. https://doi.org/10.1007/s40616-022-00179-4

Staddon, J. (2016). Learning I. In *Adaptive Behavior and Learning* (second, pp. 457–485). Cambridge University Press. https://doi.org/10.1017/CBO9781139998369.018

Staddon, J. E. R. (2016). The optimal allocation of behavior. In *Adaptive Behavior and Learning* (2nd ed., pp. 205–241). Cambridge University Press. https://doi.org/10.1017/cbo9781139998369.009

Staddon, J. E. R., & Cerutti, D. T. (2003). Operant Conditioning. *Annual Review of Psychology*, *54*, 115–144. https://doi.org/10.1146/ANNUREV.PSYCH.54.101601.145124

Staddon, J. E. R., & Staddon, J. E. R. (2016). Choice: dynamics and decision rules. In *Adaptive Behavior and Learning*. https://doi.org/10.1017/cbo9781139998369.010

Storn, R., & Price, K. (1997). Differential Evolution - A simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, *11*(4), 341–359. https://doi.org/10.1023/A:1008202821328

Stubbs, D. A., & Pliskoff, S. S. (1969). Concurrent responding with fixed relative rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, *12*(6), 887–895. https://doi.org/10.1901/JEAB.1969.12-887

Taherkhani, A., Belatreche, A., Li, Y., Cosma, G., Maguire, L. P., & McGinnity, T. M. (2020). A review of learning in biologically plausible spiking neural networks. *Neural Networks*, *122*, 253–272. https://doi.org/10.1016/j.neunet.2019.09.036

Tan, C., Šarlija, M., & Kasabov, N. (2020). Spiking neural networks: Background, recent development and the neucube architecture. *Neural Processing Letters*, *52*(2), 1675–1701. https://doi.org/10.1007/s11063-020-10322-8

Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., & Liu, C. (2018). A survey on deep transfer learning. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *11141 LNCS*, 270–279. https://doi.org/10.1007/978-3-030-01424-7_27

Tato, A., & Nkambou, R. (2022). Infusing expert knowledge into a deep neural network using attention mechanism for personalized learning environments. *Frontiers in Artificial Intelligence*, *5*(June), 1–19. https://doi.org/10.3389/frai.2022.921476

Tavanaei, A., Ghodrati, M., Kheradpisheh, S. R., Masquelier, T., & Maida, A. (2019). Deep learning in spiking neural networks. *Neural Networks*, *111*, 47–63. https://doi.org/10.1016/j.neunet.2018.12.002

Taylor, R., & Davison, M. (1983). Sensitivity to reinforcement in concurrent arithmetic and exponential schedules. *Journal of the Experimental Analysis of Behavior*, *39*(1), 191–198. https://doi.org/10.1901/jeab.1983.39-191

Taylor, S., Jaques, N., Nosakhare, E., Sano, A., & Picard, R. (2017). Personalized multitask learning for predicting tomorrow's mood, stress, and health. *IEEE Transactions on Affective Computing*, *11*(2), 200–213. https://doi.org/10.1109/TAFFC.2017.2784832

Taylor, S., Jaques, N., Nosakhare, E., Sano, A., & Picard, R. (2020). Personalized multitask learning for predicting tomorrow's mood, stress, and health. *IEEE Transactions on Affective Computing*, *11*(2), 200–213. https://doi.org/10.1109/TAFFC.2017.2784832

The Experimental Analysis of Behavior | American Scientist. (n.d.). Retrieved January 12, 2023, from https://www.americanscientist.org/article/the-experimental-analysis-of-behavior

Theobald, O. (2017). *Machine learning for absolute beginners* (2nd ed.). Theobald O.

Thompson, R. H., & Iwata, B. A. (2007). A comparison of outcomes from descriptive and functional analyses of problem behavior. *Journal of Applied Behavior Analysis*, *40*(2), 333. https://doi.org/10.1901/JABA.2007.56-06

Thorndike, E. L. (1911). *Animal Intelligence: Experimental Studies [1911]*. Transaction Publishers.

Torday, J. S., & Baluška, F. (2019). Why control an experiment?: From empiricism, via consciousness, toward Implicate Order. *EMBO Reports*, *20*(10). https://doi.org/10.15252/EMBR.201949110

Tovar, Á. E., Torres-Chávez, Á., Mofrad, A. A., & Arntzen, E. (2023). Computational models of stimulus equivalence: An intersection for the study of symbolic behavior. *Journal of the Experimental Analysis of Behavior*. https://doi.org/10.1002/JEAB.829

Tsiakmaki, M., Kostopoulos, G., Kotsiantis, S., & Ragos, O. (2020). Transfer learning from deep neural networks for predicting student performance. *Applied Sciences 2020, Vol. 10, Page 2145*, *10*(6), 2145. https://doi.org/10.3390/APP10062145

Turaev, S., Almisreb, A. A., & Saleh, M. A. (2020). Application of transfer learning for fruits and vegetable quality assessment. *Proceedings of the 2020 14th International Conference on Innovations in Information Technology, IIT 2020*, 7–12. https://doi.org/10.1109/IIT50501.2020.9299048

Turgeon, S., & Lanovaz, M. J. (2020). Tutorial: Applying machine learning in behavioral research. *Perspectives on Behavior Science*, *43*(4), 697–723. https://doi.org/10.1007/s40614-020-00270-y

Turgeon, S., & Lanovaz, M. J. (2021). Perceptions of behavior analysis in France: Accuracy and tone of posts in an internet forum on autism. *Behavior and Social Issues*, *July*. https://doi.org/10.1007/s42822-021-00057-z

Vabalas, A., Gowen, E., Poliakoff, E., & Casson, A. J. (2019). Machine learning algorithm validation with a limited sample size. *PLoS ONE*, *14*(11), 1–20. https://doi.org/10.1371/journal.pone.0224365

Vaidya, C. J., You, X., Mostofsky, S., Pereira, F., Berl, M. M., & Kenworthy, L. (2020). Data-driven identification of subtypes of executive function across typical development, attention deficit hyperactivity disorder, and autism spectrum disorders. *Journal of Child Psychology and Psychiatry and Allied Disciplines*, *61*(1), 51–61. https://doi.org/10.1111/jcpp.13114

Valletta, J. J., Torney, C., Kings, M., Thornton, A., & Madden, J. (2017). Applications of machine learning in animal behaviour studies. *Animal Behaviour*, *124*(January 2018), 203–220. https://doi.org/10.1016/j.anbehav.2016.12.005

van Gerven, M., & Bohte, S. (2017). Editorial: Artificial neural networks as models of neural information processing. *Frontiers in Computational Neuroscience*, *11*, 114. https://doi.org/10.3389/FNCOM.2017.00114

Vazquez, R. A., & Cachón, A. (2010). Integrate-and-Fire neurons and their application in pattern recognition. *Program and Abstract Book - 2010 7th International Conference on Electrical Engineering, Computing Science and Automatic Control, CCE 2010*, *Cce*, 424–428. https://doi.org/10.1109/ICEEE.2010.5608622

Verma, V. K., & Verma, S. (2022). Machine learning applications in healthcare sector: An overview. *Materials Today: Proceedings*, *57*, 2144–2147. https://doi.org/10.1016/J.MATPR.2021.12.101

Volovici, V., Syn, N. L., Ercole, A., Zhao, J. J., & Liu, N. (2022). Steps to avoid overuse and misuse of machine learning in clinical research. *Nature Medicine 2022 28:10*, *28*(10), 1996–1999. https://doi.org/10.1038/s41591-022-01961-6

Vyse, S. (2013). Changing course. *The Behavior Analyst*, *1*(1), 129–165. https://doi.org/10.9774/gleaf.978-1-907643-44-6_19

Watson, J. B. (1914). *Behavior: An introduction to comparative psychology*. Henry Holt and Company.

Weerasinghe, M. M., Ramos, J. I. E., Wang, G. Y., & Parry, D. (2021). Incorporating structural plasticity approaches in spiking neural networks for EEG modelling. *IEEE Access*, *9*, 117338–117348. https://doi.org/10.1109/ACCESS.2021.3099492

Weiss, K., Khoshgoftaar, T. M., & Wang, D. (2016). A survey of transfer learning. *Journal of Big Data 2016 3:1*, *3*(1), 1–40. https://doi.org/10.1186/S40537-016-0043-6

Wu, X., Feng, Y., Lou, S., Zheng, H., Hu, B., Hong, Z., & Tan, J. (2023). Improving NeuCube spiking neural network for EEG-based pattern recognition using transfer learning. *Neurocomputing*, *529*, 222–235. https://doi.org/10.1016/J.NEUCOM.2023.01.087

Wu, X., Kumar, V., Ross, Q. J., Ghosh, J., Yang, Q., Motoda, H., McLachlan, G. J., Ng, A., Liu, B., Yu, P. S., Zhou, Z. H., Steinbach, M., Hand, D. J., & Steinberg, D. (2008). Top 10 algorithms in data mining. *Knowledge and Information Systems*, *14*(1), 1–37. https://doi.org/10.1007/S10115-007-0114-2/METRICS

Xing, W., Liu, Y., Xin, B., Zang, L., & Deng, N. (2020). The application of deep and transfer learning for identifying cashmere and wool fibers. *Journal of Natural Fibers*, *19*(1), 88–104. https://doi.org/10.1080/15440478.2020.1727817

Yamazaki, K., Vo-Ho, V. K., Bulsara, D., & Le, N. (2022). Spiking neural networks and their applications: A review. *Brain Sciences*, *12*(7). https://doi.org/10.3390/BRAINSCI12070863

Yan, Z., Zhou, J., & Wong, W.-F. (2021). Near lossless transfer learning for spiking neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, *35*(12 SE-AAAI Technical Track on Machine Learning V), 10577–10584. https://ojs.aaai.org/index.php/AAAI/article/view/17265

Yarkoni, T., & Westfall, J. (2017). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, *12*(6), 1100–1122. https://doi.org/10.1177/1745691617693393

Young, M. E. (2018). A place for statistics in behavior analysis. *Behavior Analysis: Research and Practice*, *18*(2), 193–202. https://doi.org/10.1037/bar0000099

Zador, A. M. (2019). A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications 2019 10:1*, *10*(1), 1–7. https://doi.org/10.1038/s41467-019-11786-6

Zhan, Q., Liu, G., Xie, X., Sun, G., & Tang, H. (2021). Effective transfer learning Algorithm in spiking neural networks. *IEEE Transactions on Cybernetics*, *PP*, 1–13. https://doi.org/10.1109/TCYB.2021.3079097

Zhan, Q., Liu, G., Xie, X., Sun, G., & Tang, H. (2022). Effective transfer learning algorithm in spiking neural networks. *IEEE Transactions on Cybernetics*, *52*(12), 13323–13335. https://doi.org/10.1109/TCYB.2021.3079097

Glossary

Word/acronym/abbreviation	Meaning
Activation Function	A mathematical function that is added into an artificial neural network in order to help the network learn complex patterns in the data. When comparing with a neuron-based model that is in our brains, the activation function is at the end deciding what is to be fired to the next neuron.
Adaptive learning	Adaptive learning refers to behaviour that is acquired, maintained and changed through an organism's life.
Analogue inputs /outputs	An analog signal is one that can take on any number of values, unlike a digital signal which has only two values
Applied Behaviour Analysis (ABA)	ABA focuses on the application of EAB knowledge to real-life situations. The goal is to develop a technology of behaviour change by investigating environmental variables that reliably shape socially significant behaviour.
Artificial neural network (ANN)	ANN is a brain-inspired network model that simulates the brain in analysing and processing information.
Behaviour	This thesis refers to behaviour under the theoretical framework of behaviour analysis. Behaviour refers to anything an organism does, and it is studied in relation to environmental controlling variables. There is constant interaction between organism and environment, whereby environmental changes are followed by the organism responding to them, thereby changing the environment, and so on (see feedback function).
Behaviour Analysis	The science of psychology that studies behaviour that studies learnt behaviour in its own right, not as an index of other underlying events. There are two significant areas of behaviour analysis: experimental (EAB) and applied (ABA).
Behaviour Analyst	The psychologist that practises under the principles of behaviour analysis.
Choice and choice behaviour	It involves the distribution of responses to alternatives and the participant must choose from behaviours that exist in their own repertoire (Herrnstein, 1970). Thus, all behaviour is choice behaviour, as it involves doing one thing or the other. In EAB choice behaviour is studies extensively to understand adaptive learning. Choice is often recorded as a datum of relative rates of responding. The terms ‘choice’ and ‘choice behaviour’ are used interchangeably. **The requirements by AUT to include the Manuscripts as submitted did not allow changes.**
Class labels	Class labels in ML is the categorical or nominal values assigned to attributes whose value you want to predict based on the values of other attributes.

Word/acronym/abbreviation	Meaning
Classification	Is a predictive problem where to model prediction the correct label of a given input data. In this thesis the class labels were the components.
Component	Components are different reinforcer ratios arranged in a 2-key concurrent VI-VI schedule procedure. They are denoted as 'components' as several of those (components) change randomly within each training session. In the source experiments, seven components were arranged in each session (1:27, 1:9, 1:3, 1:1, 3:1, 9:1, 27:1), and pigeons' binary choices were observed.
Concurrent schedule procedure	A reinforcement procedure in which the participant can choose to respond on one of two or more simple reinforcement schedules that are available simultaneously. Concurrent schedules allow for the measurement of choice between simple schedule alternatives.
Cross-classification (or validation)	Cross validation is a modelling technique where a subset of samples is used to fit a model and the remaining sample is used to test the model the model's performance
Deep learning	A type of machine learning that the artificial neural networks use multiple layers to process data to extract progressively higher-level features from data. Very popular with 'big data'.
Differential equations	The occurrence of a spike is determined by differential equations that represent various "biological" processes of the machine, the most important of which is the neuron's membrane potential. They allow the estimation of the derivative of any of the network outputs with respect to any of its inputs.
Differential evolution algorithm	A popular evolutionary algorithm inspired by Darwin's theory of evolution. It has been studied extensively to solve different areas of optimisation and engineering applications since its introduction by Storn and Price (1997).
Experimental analysis of behaviour (EAB)	Experimental Analysis of Behaviour is characterised by observing measurable behaviour which could be predicted and controlled. It uses experiments to develop theories to understand behaviour-environment interactions.
F1	F1 score is a weighted average of precision and sensitivity. It is calculated to combine sensitivity and precision in one metrics.
Feedback function	A feedback function (loop) is a description of the relationship dynamics between the behaviour (response) and the environment (reinforcement). Usually in EAB it is represented as the reinforcer rate as a function of response rate.

Word/acronym/abbreviation	Meaning
Firing rate	The neuron produces a spike train (sequences of ones and zeros) from each sample (i.e., binary choices). Thus, transforms temporal data into spike trains. Then, we calculate the firing rate form the spike train. $$\mathrm{fr_i} = \frac{\mathrm{spk}}{\lvert \mathrm{s_i} \rvert}$$ The firing rate of each input pattern being spk the number of spikes.
Frequently changing environments	In choice research experiments the environmental variation is represented by multiple within session changes to the reinforcer arrangements and overall reinforcer rate of reinforcement when using concurrent schedule procedures.
Generalised Matching Law (GML)	The GML model shows how choice behaviour is governed by the relative rate of the reinforcer and is represented in the logarithmic form, $$log\left(\frac{\mathrm{B1}}{\mathrm{B2}}\right) = \mathrm{a}\, log\left(\frac{\mathrm{R1}}{\mathrm{R2}}\right) + \mathrm{logc}$$ where B is choice behaviour to alternative keys, R is reinforcer, a represents sensitivity to reinforcement, Log c represents responses for other reasons than the reinforcer.
Goal-directed behaviour	Behaviour that may be part of the individual intention to pursue a goal but which may be purely mechanical
Gradient descent	Optimisation algorithm which is commonly used to train machine learning models and neural networks
Hyperparameters	These are parameters whose values determine the learning process. They also determine the values for a learning algorithm to end up learning. The prefix 'hyper_' denote the 'top-level' parameters that control the learning process and the model parameters that result from it.
Informedness	Informedness (or Youden’s index) is the probability of an informed decision. $$\mathrm{J} = \mathrm{Sensitivity} + \mathrm{Specificity} - 1$$ $$\frac{\#\,\mathrm{TP}}{\#\,\mathrm{TP} + \mathrm{FN}} - \frac{\#\mathrm{FP}}{\mathrm{FP} + \mathrm{TN}}$$
Learning algorithms	Lists of instructions that allow neural networks to learn patterns from training data.
Learning history	The reinforcer (reinforcement) arrangement used in training conditions that led to the current choice behaviour observed by us (see also component).
Machine learning (ML)	ML is a subfield of AI those studies algorithms and techniques for automating solutions to complex problems.

Word/acronym/abbreviation	Meaning
Matching law	Mathematical equation that shows how choice is governed by the reinforcers obtained from each alternative of a concurrent schedule procedure. It has been identified as the law of simple actions.
Nested cross-validation	Often refer to as a double cross-validation process. It is a training technique capable of estimating the generalisation error of the underlying model and its hyperparameters search.
Operant chamber	A small box designed for a non-human organism that presses a lever or pecks a key when lights illuminate ***Figure 1***).
Optimisation	Is the algorithm that is used to improve the effectiveness and accuracy of a machine learning model, usually through the tweaking of model hyperparameters.
Perceptron	In a biological neuron the dendrite receives electrical signals from the axons of other neurons, in the perceptron these electrical signals are represented as numerical values. At the synapses between the dendrite and axons, electrical signals are modulated in various amounts. The neuron fires an output signal only when the total strength of the input signals exceeds a certain threshold. His output is fed to other perceptrons. In machine learning this process has inspired a single layer neural network and the perceptron consists of input values, weights and a bias, a weighted sum and activation function to model the biological neuron behaviour.
Personalised model	Personalised computational model, we use information about an individual and the available information about other individuals related to the same problem.
Precision	Precision is also called positive predictive values. $\text{Precision} = \frac{\text{True positives}}{\text{True positives} + \text{False positives}}$ See also confusion matrix
Preference pulses	Preference pulses describe the strong preference for the first few seconds to the same alternative (e.g., right-key) since the reinforcer delivery on that key (right-key). These pulses have been used to observe differences in the effects of reinforcers on behaviour. These are usually presented as logarithms of ratio (say, left: right) of responses that occur at each ordinal position after reinforcer delivery.
Reinforcement	Reinforcement is a term used in operant conditioning to refer to anything that increases the likelihood that a response will occur. Operant conditioning is a procedure where learning occurs through the association of behaviour and stimuli that is modified by reinforcement or punishment.
Reinforcement arrangement	Used interchangeably with the term reinforcer arrangement.

Word/acronym/abbreviation	Meaning
Reinforcer	A discriminable environmental stimulus that is specifically delivered following a behaviour to increase the probability that it will occur again. In applied settings many times it is used interchangeably with reward, though not accurate.
Reinforcer arrangement	Used interchangeably with the term reinforcement arrangement (also refer to learning history). **The requirements by AUT to include the Manuscripts as submitted did not allow changes.**
Relative reinforcer rate	The proportion of reinforcers obtained on one alternative of a concurrent schedule of reinforcement.
Relative response rate	The proportion of responses emitted on one alternative of a concurrent schedule of reinforcement.
Sample	A sample data set is a subset of the data. The size is always smaller than the total size of the data. Samples are fed in the machine as inputs.
Schedules of reinforcements	A set of defined rules that dictate how a stimulus can be manipulated to strengthen a specific behaviour's probability to occur in the future. In other words, the rules specify when a response will be reinforced.
Sensitivity (recall)	It measures the accuracy in the event population; the true positive rate. $\mathrm{Sensitivity} = \frac{\text{\# of samples predicted as having the event}}{\text{\# of samples with the event of interest}}$
Sensitivity (to reinforcement)	Sensitivity to reinforcement is an index to denote behavioural change in response to environmental conditions. It is calculated in the GML logarithmic form $log\left(\frac{\mathrm{B1}}{\mathrm{B2}}\right) = \mathrm{a}\, log\left(\frac{\mathrm{R1}}{\mathrm{R2}}\right) + \mathrm{logc}$ where parameter a is sensitivity to the reinforcement (the slope of the best-fitted line) and measures the degree to which the response ratio changes with variations in the reinforcer ratio.
Single-subject design (SSD)	In an SSD the subject serves as its own control, rather than using another individual/group. This is done by using repeated measurements to really understand an individual's variability
Specificity	Also (1-specificity); the false positive rate. $\mathrm{Specificity} = \frac{\text{\# of samples predicted as non events}}{\text{\# of samples without the event of interest}}$
Spikes	Discrete events placed at a point of time, rather than continuous values. Spikes are used with spiking neural networks.
Spiking neural network (SNN)	Spiking neural networks is the third generation of artificial neural networks (ANN) with nodes that mimic the dynamics of biological neurons. Unlike previous generations, this type of neural networks incorporates the concept of time and space.

Word/acronym/abbreviation	Meaning
Steady-state procedures	That meant that the reinforcer conditions remained the same with the passage of time (up to 30 daily sessions) and only changed according to set criteria by the experimenter.
The Leaky Integrate-and-Fire (LIF)	The most common computational neuron model that defines the evolution of the membrane potential but also the mechanism to generate spikes with low computational cost.
Transfer learning (TL)	TL is a ML method that uses a pre-trained model as a starting point to develop a new artificial model to improve the classification/prediction process
Variable interval (VI) schedule	When using a VI schedule the response is reinforced when a variable amount of time has elapsed since the last reinforcer
Vector	Vector is an array of numbers
Weights	Weight is a parameter used with neural networks that transforms input data within the layer (see also neural network). It represents the strength of the connection between units. They are attached with each input/feature and convey the importance of that corresponding feature in predicting the final output.

Appendices

Appendix A: Permissions

PERMISSION FOR DATASETS

A1. Permission from Associate Professor Landon to publish results generated from Landon and Davison (2001) datasets

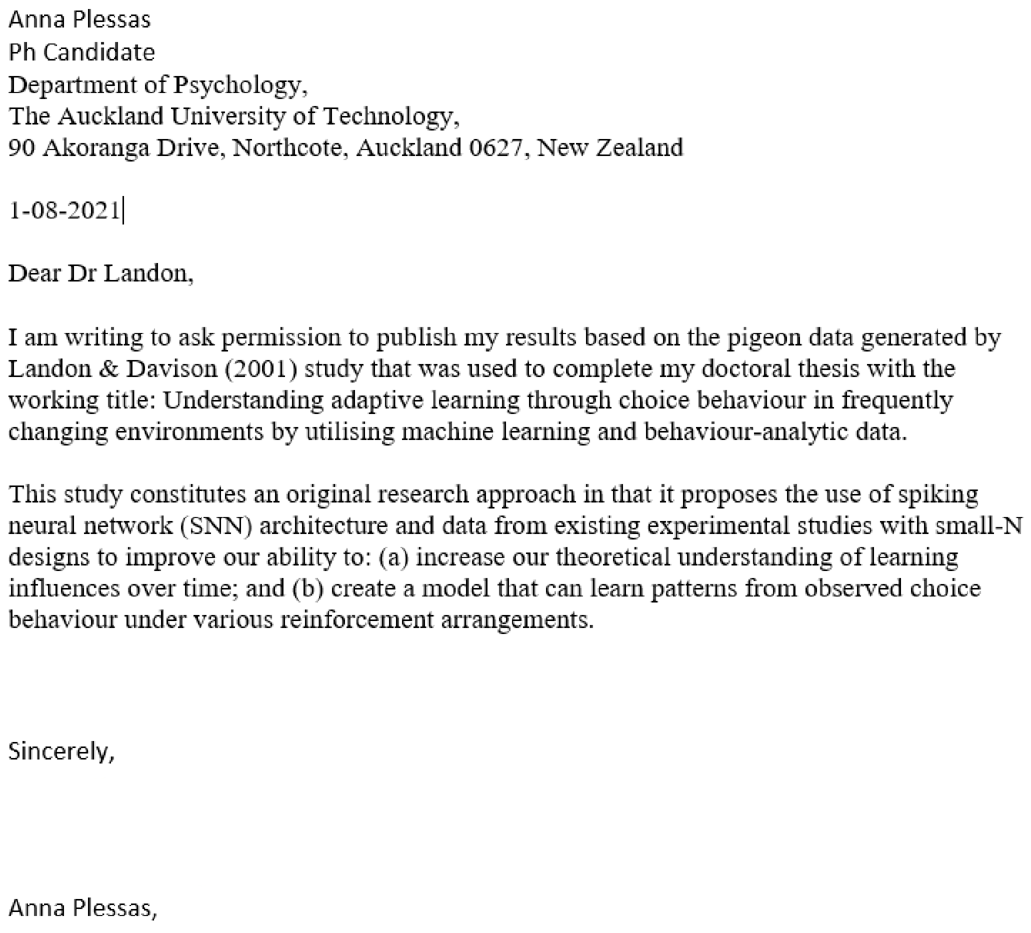

Anna Plessas
Ph Candidate
Department of Psychology,
The Auckland University of Technology,
90 Akoranga Drive, Northcote, Auckland 0627, New Zealand

1-08-2021

Dear Dr Landon,

I am writing to ask permission to publish my results based on the pigeon data generated by Landon & Davison (2001) study that was used to complete my doctoral thesis with the working title: Understanding adaptive learning through choice behaviour in frequently changing environments by utilising machine learning and behaviour-analytic data.

This study constitutes an original research approach in that it proposes the use of spiking neural network (SNN) architecture and data from existing experimental studies with small-N designs to improve our ability to: (a) increase our theoretical understanding of learning influences over time; and (b) create a model that can learn patterns from observed choice behaviour under various reinforcement arrangements.

Sincerely,

Anna Plessas,

Permission received from:

Name of author: Dr Jason Landon

Signature:

A2. Permission from Emeritus Professor Davison to publish results generated from Landon and Davison (2001) datasets

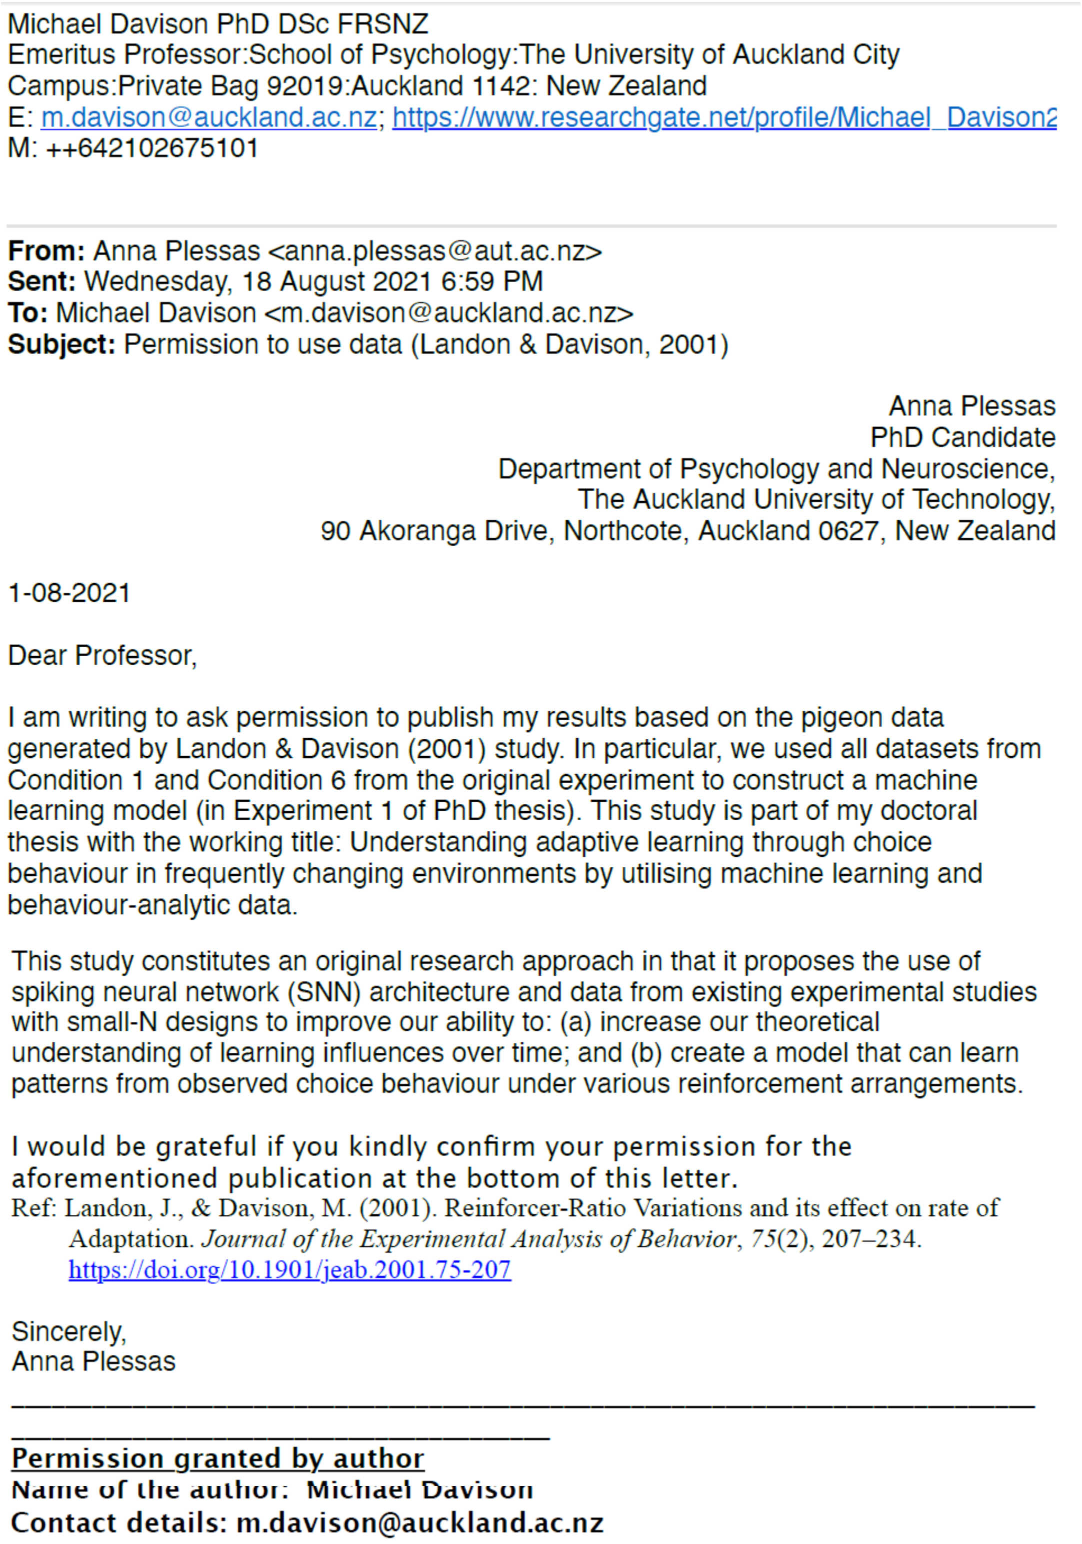

Michael Davison PhD DSc FRSNZ
Emeritus Professor:School of Psychology:The University of Auckland City Campus:Private Bag 92019:Auckland 1142: New Zealand
E: m.davison@auckland.ac.nz; https://www.researchgate.net/profile/Michael_Davison2
M: ++642102675101

From: Anna Plessas <anna.plessas@aut.ac.nz>
Sent: Wednesday, 18 August 2021 6:59 PM
To: Michael Davison <m.davison@auckland.ac.nz>
Subject: Permission to use data (Landon & Davison, 2001)

Anna Plessas
PhD Candidate
Department of Psychology and Neuroscience,
The Auckland University of Technology,
90 Akoranga Drive, Northcote, Auckland 0627, New Zealand

1-08-2021

Dear Professor,

I am writing to ask permission to publish my results based on the pigeon data generated by Landon & Davison (2001) study. In particular, we used all datasets from Condition 1 and Condition 6 from the original experiment to construct a machine learning model (in Experiment 1 of PhD thesis). This study is part of my doctoral thesis with the working title: Understanding adaptive learning through choice behaviour in frequently changing environments by utilising machine learning and behaviour-analytic data.

This study constitutes an original research approach in that it proposes the use of spiking neural network (SNN) architecture and data from existing experimental studies with small-N designs to improve our ability to: (a) increase our theoretical understanding of learning influences over time; and (b) create a model that can learn patterns from observed choice behaviour under various reinforcement arrangements.

I would be grateful if you kindly confirm your permission for the aforementioned publication at the bottom of this letter.

Ref: Landon, J., & Davison, M. (2001). Reinforcer-Ratio Variations and its effect on rate of Adaptation. *Journal of the Experimental Analysis of Behavior*, *75*(2), 207–234. https://doi.org/10.1901/jeab.2001.75-207

Sincerely,
Anna Plessas

Permission granted by author
Name of the author: Michael Davison
Contact details: m.davison@auckland.ac.nz

A3. Permission from Emeritus Professor to publish results generated from Davison & Baum (2000) datasets

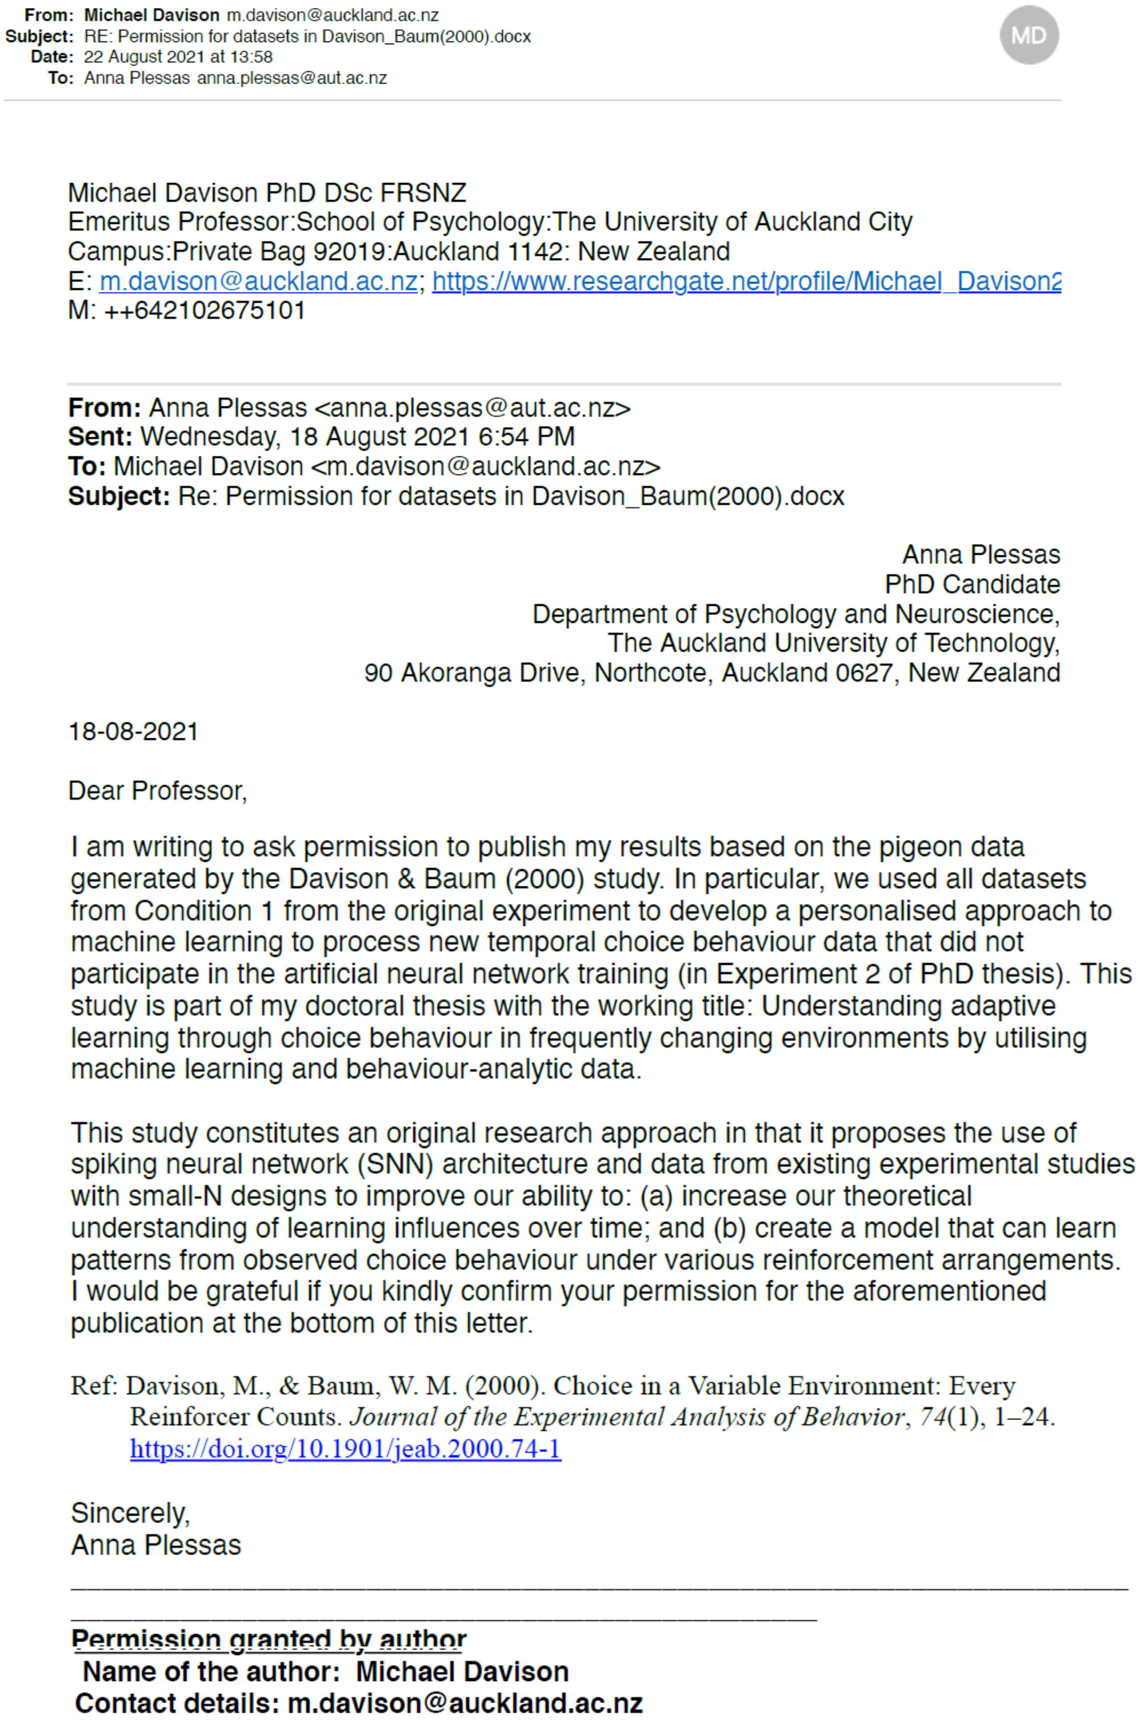

From: **Michael Davison** m.davison@auckland.ac.nz
Subject: RE: Permission for datasets in Davison_Baum(2000).docx
Date: 22 August 2021 at 13:58
To: Anna Plessas anna.plessas@aut.ac.nz

Michael Davison PhD DSc FRSNZ
Emeritus Professor:School of Psychology:The University of Auckland City Campus:Private Bag 92019:Auckland 1142: New Zealand
E: m.davison@auckland.ac.nz; https://www.researchgate.net/profile/Michael_Davison2
M: ++642102675101

From: Anna Plessas <anna.plessas@aut.ac.nz>
Sent: Wednesday, 18 August 2021 6:54 PM
To: Michael Davison <m.davison@auckland.ac.nz>
Subject: Re: Permission for datasets in Davison_Baum(2000).docx

Anna Plessas
PhD Candidate
Department of Psychology and Neuroscience,
The Auckland University of Technology,
90 Akoranga Drive, Northcote, Auckland 0627, New Zealand

18-08-2021

Dear Professor,

I am writing to ask permission to publish my results based on the pigeon data generated by the Davison & Baum (2000) study. In particular, we used all datasets from Condition 1 from the original experiment to develop a personalised approach to machine learning to process new temporal choice behaviour data that did not participate in the artificial neural network training (in Experiment 2 of PhD thesis). This study is part of my doctoral thesis with the working title: Understanding adaptive learning through choice behaviour in frequently changing environments by utilising machine learning and behaviour-analytic data.

This study constitutes an original research approach in that it proposes the use of spiking neural network (SNN) architecture and data from existing experimental studies with small-N designs to improve our ability to: (a) increase our theoretical understanding of learning influences over time; and (b) create a model that can learn patterns from observed choice behaviour under various reinforcement arrangements. I would be grateful if you kindly confirm your permission for the aforementioned publication at the bottom of this letter.

Ref: Davison, M., & Baum, W. M. (2000). Choice in a Variable Environment: Every Reinforcer Counts. *Journal of the Experimental Analysis of Behavior*, *74*(1), 1–24. https://doi.org/10.1901/jeab.2000.74-1

Sincerely,
Anna Plessas

Permission granted by author
Name of the author: Michael Davison
Contact details: m.davison@auckland.ac.nz

A4. Permission from Emeritus Professor Baum to publish results generated from Davison and Baum (2000) datasets

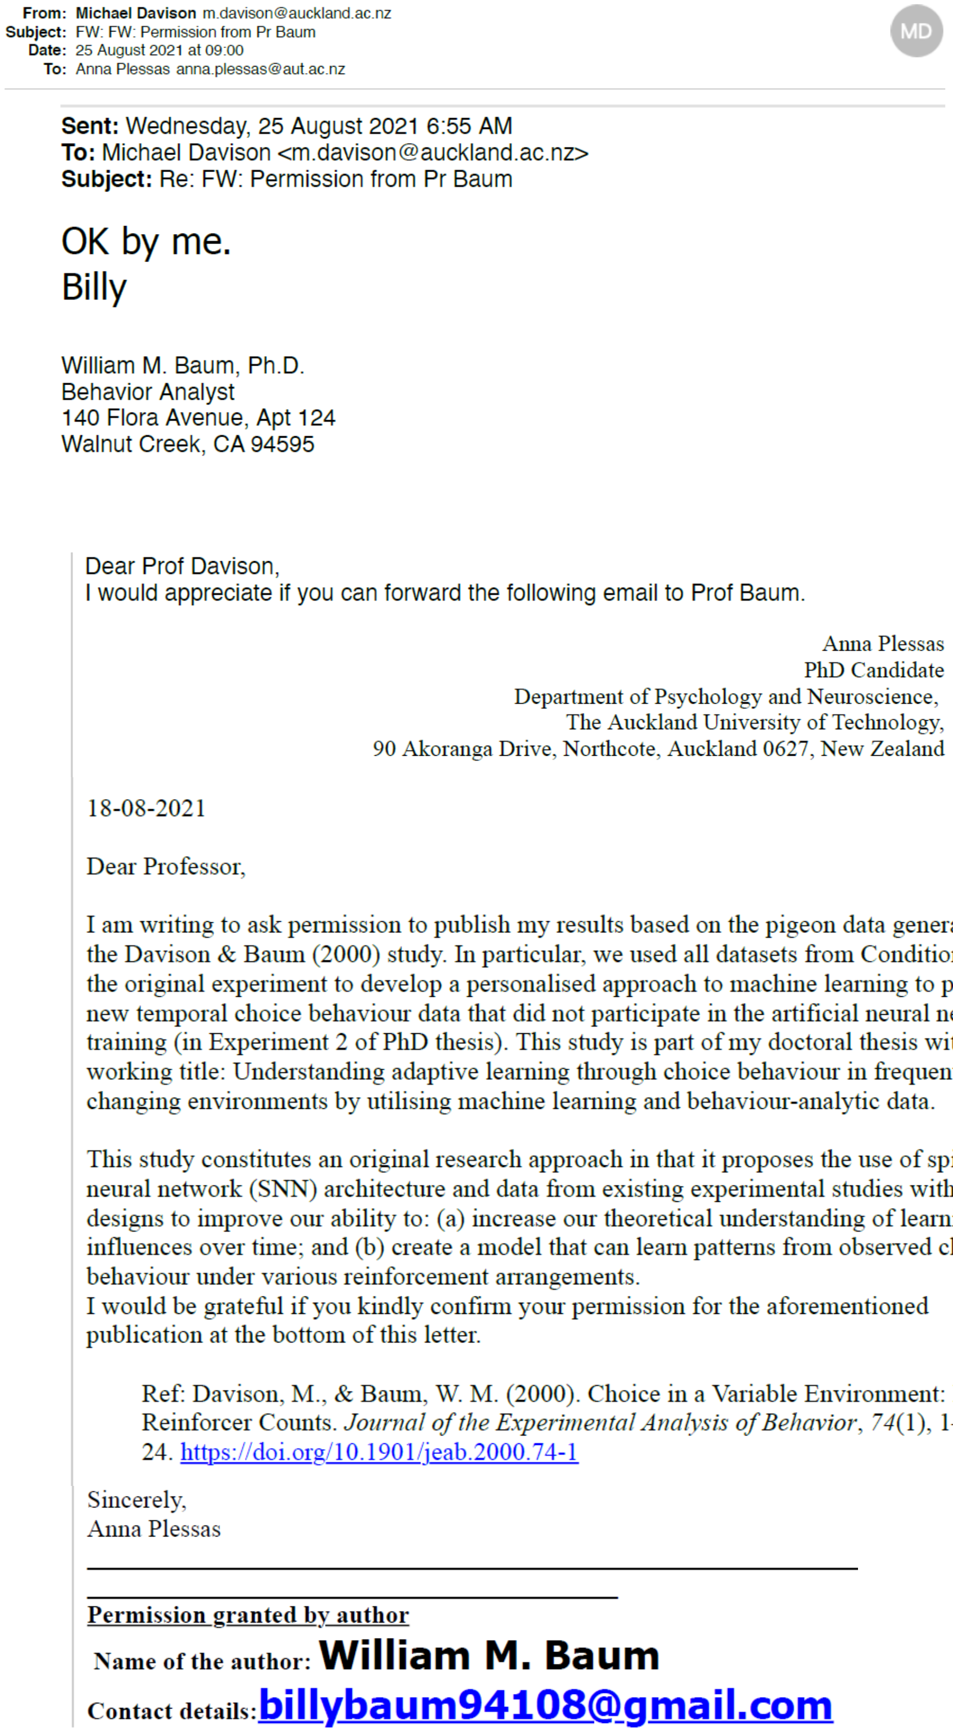

From: **Michael Davison** m.davison@auckland.ac.nz
Subject: FW: FW: Permission from Pr Baum
Date: 25 August 2021 at 09:00
To: Anna Plessas anna.plessas@aut.ac.nz

Sent: Wednesday, 25 August 2021 6:55 AM
To: Michael Davison <m.davison@auckland.ac.nz>
Subject: Re: FW: Permission from Pr Baum

OK by me.
Billy

William M. Baum, Ph.D.
Behavior Analyst
140 Flora Avenue, Apt 124
Walnut Creek, CA 94595

> Dear Prof Davison,
> I would appreciate if you can forward the following email to Prof Baum.
>
> Anna Plessas
> PhD Candidate
> Department of Psychology and Neuroscience,
> The Auckland University of Technology,
> 90 Akoranga Drive, Northcote, Auckland 0627, New Zealand
>
> 18-08-2021
>
> Dear Professor,
>
> I am writing to ask permission to publish my results based on the pigeon data generated by the Davison & Baum (2000) study. In particular, we used all datasets from Condition 1 from the original experiment to develop a personalised approach to machine learning to process new temporal choice behaviour data that did not participate in the artificial neural network training (in Experiment 2 of PhD thesis). This study is part of my doctoral thesis with the working title: Understanding adaptive learning through choice behaviour in frequently changing environments by utilising machine learning and behaviour-analytic data.
>
> This study constitutes an original research approach in that it proposes the use of spiking neural network (SNN) architecture and data from existing experimental studies with small-N designs to improve our ability to: (a) increase our theoretical understanding of learning influences over time; and (b) create a model that can learn patterns from observed choice behaviour under various reinforcement arrangements.
> I would be grateful if you kindly confirm your permission for the aforementioned publication at the bottom of this letter.
>
> Ref: Davison, M., & Baum, W. M. (2000). Choice in a Variable Environment: Every Reinforcer Counts. *Journal of the Experimental Analysis of Behavior*, *74*(1), 1–24. https://doi.org/10.1901/jeab.2000.74-1
>
> Sincerely,
> Anna Plessas
>
> **Permission granted by author**
> **Name of the author:** **William M. Baum**
> **Contact details:** **billybaum94108@gmail.com**

A5. Permission from Associate Professor Landon to publish results generated from Landon et. al (2007) datasets

Anna Plessas
PhD Candidate
Department of Psychology and Neuroscience,
The Auckland University of Technology,
90 Akoranga Drive, Northcote, Auckland 0627, New Zealand

1-08-2021

Dear Professor,

I am writing to ask permission to publish my results based on the pigeon data generated by Landon et al. (200) study. In particular, we used all datasets from Condition 1 from the original experiment to construct a machine learning model (in Experiment 3 of PhD thesis). This study is part of my doctoral thesis with the working title: Understanding adaptive learning through choice behaviour in frequently changing environments by utilising machine learning and behaviour-analytic data.

This study constitutes an original research approach in that it proposes the use of spiking neural network (SNN) architecture and data from existing experimental studies with small-N designs to improve our ability to: (a) increase our theoretical understanding of learning influences over time; and (b) create a model that can learn patterns from observed choice behaviour under various reinforcement arrangements.

> Ref: Landon, J; Davison, M; Krägeloh, CU; Thompson, NM; Miles, JL; Vickers, MH; Fraser, M and Breier, B. (2007). Global undernutrition during gestation. Learning & Behavior, 35(2), 79–86.

Sincerely,

Anna Plessas

Permission granted by author

Name of the author: Dr Jason Landon

Signature:

Contact details: jlandon@aut.ac.nz

A6. Permission from Associate Professor Krägeloh to publish results generated from Krägeloh et a. (2010) datasets

Anna Plessas
PhD Candidate
Department of Psychology and Neuroscience,
Auckland University of Technology,
90 Akoranga Drive, Northcote, Auckland 0627, New Zealand

1-08-2021

Dear Dr Krägeloh,

I am writing to ask permission to publish my results based on the human data generated by Krägeloh et al. (2010) study that was used to complete the third experiment of my doctoral thesis with the working title: Understanding adaptive learning through choice behaviour in frequently changing environments by utilising machine learning and behaviour-analytic data.

This study constitutes an original research approach in that it proposes the use of spiking neural network (SNN) architecture and data from existing experimental studies with small-N designs to improve our ability to: (a) increase our theoretical understanding of learning influences over time; and (b) create a model that can learn patterns from observed choice behaviour under various reinforcement arrangements.

Sincerely,

Anna Plessas,

Permission received by:

Name of the author: Chris Krägeloh

Signature:

Contact details: chris.krageloh@aut.ac.nz

AGREEMENTS FOR DISCLOSURE

A7. Nondisclosure Agreement between KEDRI and Anna Plessas

AUT KNOWLEDGE ENGINEERING & DISCOVERY RESEARCH INSTITUTE

Prof Nikola Kasabov, PhD, FIEEE, FRSNZ, FINNS College of Fellows, DVF RAE UK
Founding Director, Knowledge Engineering & Discovery Research Institute
Auckland University of Technology (AUT), Auckland, New Zealand

This Nondisclosure Agreement (the "Agreement") is entered into by and between The Knowledge Discovery & Research Institute (KEDRI) and Anna Plessas, preventing the unauthorised disclosure of Confidential Information pertaining to the use of NeuCube and its source code.

Anna Plessas agrees to hold and maintain the Confidential Information in strictest confidence for the sole and exclusive benefit of KEDRI and will restrict access to Confidential Information to third parties. She will not benefit, publish, copy, or otherwise disclose to others any information deemed confidential. Any changes made to the source code is to be informed directly to the Director of KEDRI, Prof. Nikola Kasabov. Anna Plessas will return to KEDRI any and all records, notes, and other written, printed, or tangible materials in her possession pertaining to Confidential Information immediately, if KEDRI requests it in writing.

Prof. Nikola Kasabov, FIEEE, FRSNZ, FIITP NZ, DVF RAE and SICSA UK
Director KEDRI – Knowledge Engineering and Discovery Research Institute, www.kedri.aut.ac.nz
Professor, School of Engineering, Computer and Mathematical Sciences, Auckland University of Technology
Advisory Professor Shanghai Jiao Tong University; Visiting Professor RGU UK.
President Asia-Pacific Neural Network Society (APNNS) 2019.
Honourable Member of Bulgarian Academic Soc. For Comp.Sys. and IT (AOKSIT) and Greek Computer Society.
Profile: https://scholar.google.com/citations?user=YTa9Dz4AAAAJ
Editor: Springer Handbook of Bio-/Neuroinformatics: http://www.springer.com/gp/book/9783642305733
Editor: Springer Series on Bio- and Neurosystems: http://www.springer.com/series/15821
Co-EIC: Springer journal *Evolving Systems:* https://link.springer.com/journal/12530
Science Co-Ordinator, PANTHER PROJECT
P +64 9 921 9506 **M** +64 21 488 214 **E** nkasabov@aut.ac.nz **W** www.kedri.aut.ac.nz
Coordinator of the PANTHER programme for New Zealand

Dated: 16.4.20

Anna Plessas
Dated:16.4.20

AUT Tower, Level 7, 2-14 Wakefield Street, Auckland 1010, Private Bag 92006, Auckland 1142, New Zealand.
P: 64 9 9219506 E: nkasabov@aut.ac.nz W: www.kedri.aut.ac.nz

A8. Permission request to Wiley Online Library for use of copyright manuscript for this thesis and the agreement sent from Wiley Online Library

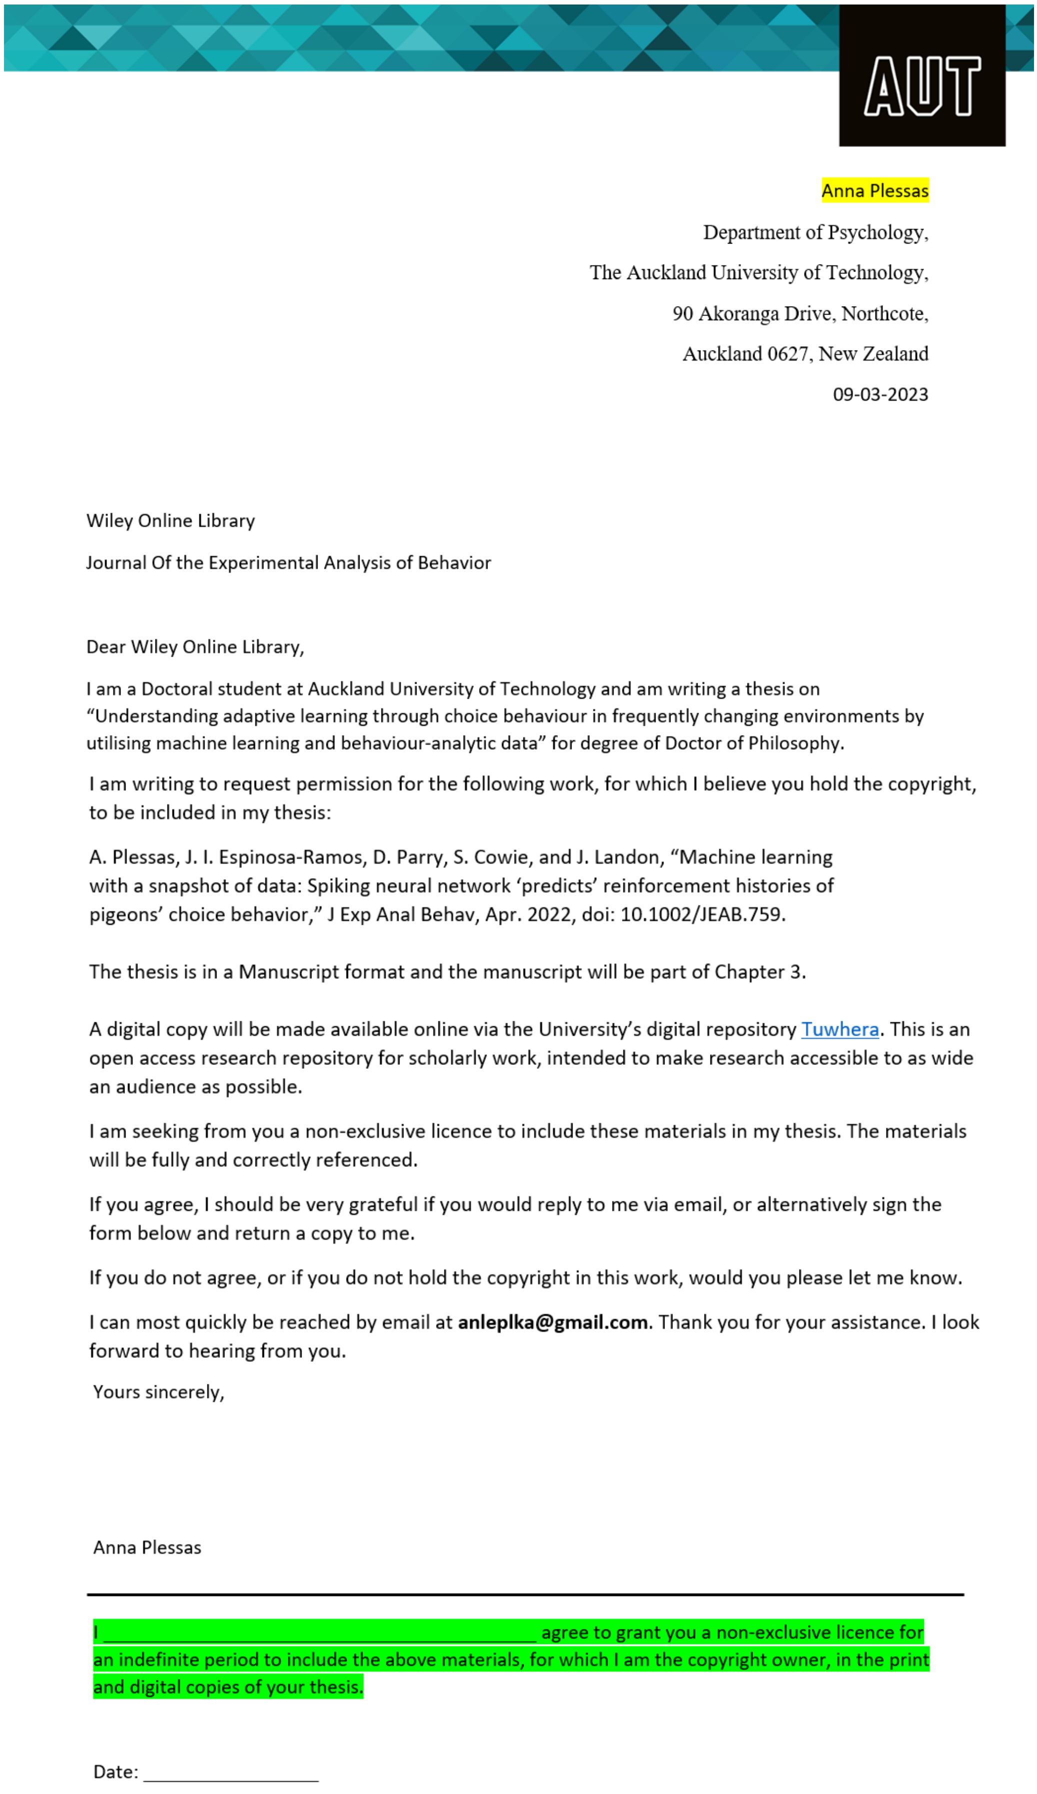

Anna Plessas
Department of Psychology,
The Auckland University of Technology,
90 Akoranga Drive, Northcote,
Auckland 0627, New Zealand

09-03-2023

Wiley Online Library

Journal Of the Experimental Analysis of Behavior

Dear Wiley Online Library,

I am a Doctoral student at Auckland University of Technology and am writing a thesis on "Understanding adaptive learning through choice behaviour in frequently changing environments by utilising machine learning and behaviour-analytic data" for degree of Doctor of Philosophy.

I am writing to request permission for the following work, for which I believe you hold the copyright, to be included in my thesis:

A. Plessas, J. I. Espinosa-Ramos, D. Parry, S. Cowie, and J. Landon, "Machine learning with a snapshot of data: Spiking neural network 'predicts' reinforcement histories of pigeons' choice behavior," J Exp Anal Behav, Apr. 2022, doi: 10.1002/JEAB.759.

The thesis is in a Manuscript format and the manuscript will be part of Chapter 3.

A digital copy will be made available online via the University's digital repository Tuwhera. This is an open access research repository for scholarly work, intended to make research accessible to as wide an audience as possible.

I am seeking from you a non-exclusive licence to include these materials in my thesis. The materials will be fully and correctly referenced.

If you agree, I should be very grateful if you would reply to me via email, or alternatively sign the form below and return a copy to me.

If you do not agree, or if you do not hold the copyright in this work, would you please let me know.

I can most quickly be reached by email at **anleplka@gmail.com**. Thank you for your assistance. I look forward to hearing from you.

Yours sincerely,

Anna Plessas

I _______________________________________ agree to grant you a non-exclusive licence for an indefinite period to include the above materials, for which I am the copyright owner, in the print and digital copies of your thesis.

Date: _______________

JOHN WILEY AND SONS LICENSE TERMS AND CONDITIONS
Mar 09, 2023

This Agreement between Mrs. Anna Plessas ("You") and John Wiley and Sons ("John Wiley and Sons") consists of your license details and the terms and conditions provided by John Wiley and Sons and Copyright Clearance Center.
License Number 5504680702368
License date Mar 09, 2023
Licensed Content Publisher

John Wiley and Sons
Licensed Content Publication

Journal of the Experimental Analysis of Behavior
Licensed
Machine learning with a snapshot of data: Spiking neural network ‘predicts’
Content Title reinforcement histories of pigeons' choice behavior
Licensed Content Author

Jason Landon, Sarah Cowie, Dave Parry, et al

Licensed Content Date
Apr 21, 2022

Licensed Content Volume 117

Licensed Content Issue 3

Licensed Content Pages 19
Type of use Dissertation/Thesis

Requestor type Author of this Wiley article

Format Electronic
Portion Full article
Will you be translating? No

Title: Understanding adaptive learning through choice behaviour in frequently changing environments by utilising machine learning and behaviour-analytic data

Institution name Auckland University Of Technology

Expected presentation date Apr 2023

Order reference number
Machine learning with a snapshot of data: Spiking neural network ‘predicts’ reinforcement histories of pigeons' choice behavior

Requestor Location
Mrs. Anna Plessas 1 Jojos way Lynmore
BOP
Rotorua, 3076 New Zealand
Attn: Mrs. Anna Plessas

Publisher Tax EU826007151 ID
Total 0.00 USD Terms and Conditions

TERMS AND CONDITIONS

This copyrighted material is owned by or exclusively licensed to John Wiley & Sons, Inc. or one of its group companies (each a"Wiley Company") or handled on behalf of a society with which a Wiley Company has exclusive publishing rights in relation to a particular work (collectively "WILEY"). By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the billing and payment terms and conditions established by the Copyright Clearance Center Inc., ("CCC's Billing and Payment terms and conditions"), at the time that you opened your RightsLink account (these are available at any time at http://myaccount.copyright.com).

Terms and Conditions

The materials you have requested permission to reproduce or reuse (the "Wiley Materials") are protected by copyright.

You are hereby granted a personal, non-exclusive, non-sub licensable (on a stand- alone basis), non-transferable, worldwide, limited license to reproduce the Wiley Materials for the purpose specified in the licensing process. This license, and any CONTENT (PDF or image file) purchased as part of your order, is for a one-time use only and limited to any maximum distribution number specified in the license. The first instance of republication or reuse granted by this license must be completed within two years of the date of the grant of this license (although copies prepared before the end date may be distributed thereafter). The Wiley Materials shall not be used in any other manner or for any other purpose, beyond what is granted in the license. Permission is granted subject to an appropriate acknowledgement given to the author, title of the material/book/journal and the publisher. You shall also duplicate the copyright notice that appears in the Wiley publication in your use of the Wiley Material. Permission is also granted on the understanding that nowhere in the text is a previously published source acknowledged for all or part of this Wiley Material. Any third party content is expressly excluded from this permission.

With respect to the Wiley Materials, all rights are reserved. Except as expressly granted by the terms of the license, no part of the Wiley Materials may be copied, modified, adapted (except for minor reformatting required by the new Publication), translated, reproduced, transferred or distributed, in any form or by any means, and no derivative works may be made based on the Wiley Materials without the prior permission of the respective copyright owner.**For STM Signatory Publishers clearing permission under the terms of the STM Permissions Guidelines only, the terms of the license are extended to include subsequent editions and for editions in other languages, provided such editions are for the work as a whole in situ and does not involve the separate exploitation of the permitted figures or extracts,** You may not alter, remove or suppress in any manner any copyright, trademark or other notices displayed by the Wiley Materials. You may not license, rent, sell, loan, lease, pledge, offer as security, transfer or assign the Wiley Materials on a stand-alone basis, or any of the rights granted to you hereunder to any other person.

The Wiley Materials and all of the intellectual property rights therein shall at all times remain the exclusive property of John Wiley & Sons Inc, the Wiley Companies, or their respective licensors, and your interest therein is only that of having possession of and the right to reproduce the Wiley Materials pursuant to Section 2 herein during the continuance of this Agreement. You agree that you own no right, title or interest in or to the Wiley Materials or any of the intellectual property rights therein. You shall have no rights hereunder other than the license as provided for above in Section 2. No right, license or interest to any trademark, trade name, service mark or other branding ("Marks") of WILEY or its licensors is granted hereunder, and you agree that you shall not assert any such right, license or interest with respect thereto

NEITHER WILEY NOR ITS LICENSORS MAKES ANY WARRANTY OR REPRESENTATION OF ANY KIND TO YOU OR ANY THIRD PARTY, EXPRESS, IMPLIED OR STATUTORY, WITH RESPECT TO THE MATERIALS OR THE ACCURACY OF ANY INFORMATION CONTAINED IN THE MATERIALS, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTY OF MERCHANTABILITY, ACCURACY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, USABILITY, INTEGRATION OR NON-INFRINGEMENT AND ALL SUCH WARRANTIES ARE HEREBY EXCLUDED BY WILEY AND ITS LICENSORS AND WAIVED BY YOU.

WILEY shall have the right to terminate this Agreement immediately upon breach of this Agreement by you.

You shall indemnify, defend and hold harmless WILEY, its Licensors and their respective directors, officers, agents and employees, from and against any actual or threatened claims, demands, causes of action or proceedings arising from any breach of this Agreement by you.

IN NO EVENT SHALL WILEY OR ITS LICENSORS BE LIABLE TO YOU OR ANY OTHER PARTY OR ANY OTHER PERSON OR ENTITY FOR ANY SPECIAL, CONSEQUENTIAL, INCIDENTAL, INDIRECT, EXEMPLARY OR PUNITIVE DAMAGES, HOWEVER CAUSED, ARISING OUT OF OR IN CONNECTION WITH THE DOWNLOADING, PROVISIONING, VIEWING OR USE OF THE MATERIALS REGARDLESS OF THE FORM OF ACTION, WHETHER FOR BREACH OF CONTRACT, BREACH OF WARRANTY, TORT, NEGLIGENCE, INFRINGEMENT OR OTHERWISE (INCLUDING, WITHOUT LIMITATION, DAMAGES BASED ON LOSS OF PROFITS, DATA, FILES, USE, BUSINESS OPPORTUNITY OR CLAIMS OF THIRD PARTIES), AND WHETHER OR NOT THE PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY PROVIDED HEREIN.

Should any provision of this Agreement be held by a court of competent jurisdiction to be illegal, invalid, or unenforceable, that provision shall be deemed amended to achieve as nearly as possible the same economic effect as the original provision, and the legality, validity and enforceability of the remaining provisions of this Agreement shall not be affected or impaired thereby.

The failure of either party to enforce any term or condition of this Agreement shall not constitute a waiver of either party's right to enforce each and every term and condition of this Agreement. No breach under this agreement shall be deemed waived or excused by either party unless such waiver or consent is in writing signed by the party granting such waiver or consent. The waiver by or consent of a party to a breach of any provision of this Agreement shall not operate or be construed as a waiver of or consent to any other or subsequent breach by such other party.

This Agreement may not be assigned (including by operation of law or otherwise) by you without WILEY's prior written consent.

Any fee required for this permission shall be non-refundable after thirty (30) days from receipt by the CCC.

These terms and conditions together with CCC's Billing and Payment terms and conditions (which are incorporated herein) form the entire agreement between you and WILEY concerning this licensing transaction and (in the absence of fraud) supersedes all prior agreements and representations of the parties, oral or written. This Agreement may not be amended except in writing signed by both parties. This Agreement shall be binding upon and inure to the benefit of the parties' successors, legal representatives, and authorized assigns.

In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall prevail.

WILEY expressly reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.

This Agreement will be void if the Type of Use, Format, Circulation, or Requestor Type was misrepresented during the licensing process.

This Agreement shall be governed by and construed in accordance with the laws of the State of New York, USA, without regards to such state's conflict of law rules. Any legal action, suit or proceeding arising out of or relating to these Terms and Conditions or the breach thereof shall be instituted in a court of competent jurisdiction in New York County in the State of New York in the United States of America and each party hereby consents and submits to the personal jurisdiction of such court, waives any objection to venue in such court and consents to service of process by registered or certified mail, return receipt requested, at the last known address of such party.

WILEY OPEN ACCESS TERMS AND CONDITIONS

Wiley Publishes Open Access Articles in fully Open Access Journals and in Subscription journals offering Online Open. Although most of the fully Open Access journals publish open access articles under the terms of the Creative Commons Attribution (CC BY) License only, the subscription journals and a few of the Open Access Journals offer a choice of Creative Commons Licenses. The license type is clearly identified on the article.

The Creative Commons Attribution License

The Creative Commons Attribution License (CC-BY) allows users to copy, distribute and transmit an article, adapt the article and make commercial use of the article. The CC-BY license permits commercial and non-

Creative Commons Attribution Non-Commercial License

The Creative Commons Attribution Non-Commercial (CC-BY-NC)License permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.(see below)

Creative Commons Attribution-Non-Commercial-NoDerivs License

The Creative Commons Attribution Non-Commercial-NoDerivs License (CC-BY-NC-ND) permits use, distribution and reproduction in any medium, provided the original work is properly cited, is not used for commercial purposes and no modifications or adaptations are made. (see below)

Use by commercial "for-profit" organizations

Use of Wiley Open Access articles for commercial, promotional, or marketing purposes requires further explicit permission from Wiley and will be subject to a fee.

Further details can be found on Wiley Online Library http://olabout.wiley.com/WileyCDA/Section/id-410895.html

Other Terms and Conditions:

v1.10 Last updated September 2015 Questions? customercare@copyright.com.

Appendix B: Chapter 4 Supplementary Material

B1. Experimenting with Dataset Approaches

Each data point was approached in two ways, by extracting information based on the time and on the actual location in which the reinforcer was delivered in relation to pigeons' choices (Table 1). By using this method, the data were amenable to analysis by SNN and we were able to provide comparisons of the performance in relation to the measure chosen to present the data. No feature extraction techniques were applied, as no irrelevant data were involved in the process.

Table 1

The two approaches used to measure the temporal pigeon data within samples

#	Description	Ratio	Example
1	Responses that occurred within 5 sec of the reinforcer being allocated with p=0.037 AND dependent on the reinforcer location	L/R or R/L	If a reinforcer was allocated to the R and, after consuming the reinforcer, the pigeon started pecking 1 time on the L and then 4 times on the R, the relative response of the pigeon will be 4/1(the pigeon pecked 4 more times on the R). If the reinforcer was allocated to the L and the pigeon pecked 1 time on the R and 3 times on the L, then the relative response will be 3/1 (the pigeon pecked 3 more times on the L)
2	Responses that occurred within 5 sec of the reinforcer being allocated with p=0.037 AND ratios of L or R responses to total responses	L/L+R Or R/R+L	If a reinforcer was allocated to the R and, after consuming the reinforcer, the pigeon started pecking 1 time on the L and then 4 times on the R, the relative response of the pigeon will be 4/5. If the reinforcer was located to the L and the pigeon pecked 1 time on the R and 3 times on the L, then the relative response will be 3/4

After conducting an analysis using both approaches, approach #2 was decided that cannot capture the relative responses to location which reflect the reinforcement arrangement and reinforcer provided. It also failed to capture the relative responses aiming for and produced less models when trained the SNN (Table 2).

Table 2

Total number of successful models out of ten runs, using pigeon data from five pigeons for training and the data of the remaining pigeon for testing, for both approaches

	Pigeon 61	Pigeon 62	Pigeon 63	Pigeon 64	Pigeon 65	Pigeon 66	Total
A1	10	10	10	9	10	10	59/60 (98.3%)
A2	5	6	9	7	5	8	42/60 (66.66%)

*Note: the columns' headings denote which pigeon's data was used for testing

Subsequently, the SNN model chosen was the one with the highest overall metrics under each combination of pigeons, for both approaches. An average value from training and testing accuracy was calculated for each in order to evaluate performances (Table 3). As we can see, all models produced a high overall accuracy with both approaches. With both approaches, the lowest result was when bird-66 data were used for testing (93%) and the highest with birds 61, 63 and 64 (96%). Despite high accuracy, further analysis was conducted to determine types of errors and how these might have affected the model's performance.

Table 3

The overall results of the best model for both approaches

#		Pigeon 61	Pigeon 62	Pigeon 63	Pigeon 64	Pigeon 65	Pigeon 66
		overall	overall	overall	overall	overall	overall
A1	Accuracy	0.96	0.96	0.96	0.96	0.94	0.93
	Sensitivity	0.87	0.87	0.87	0.86	0.81	0.81
	Specificity	0.98	0.98	0.97	0.97	0.96	0.96
	Informedness	0.85	0.85	0.84	0.83	0.77	0.77
	Precision	0.87	0.87	0.87	0.86	0.81	0.81
	F1	0.87	0.87	0.87	0.86	0.81	0.81
A2	Accuracy	0.96	0.95	0.96	0.96	0.95	0.93
	Sensitivity	0.89	0.85	0.88	0.87	0.84	0.79
	Specificity	0.98	0.97	0.98	0.98	0.97	0.96
	Informedness	0.87	0.82	0.86	0.85	0.81	0.75
	Precision	0.89	0.85	0.88	0.87	0.84	0.79
	F1	0.89	0.85	0.88	0.87	0.84	0.79

*Note: the results reflect cross-validation (CV) training and testing for validation. The results are listed per pigeon left out for testing.

Further investigation was conducted to investigate if the predictive ability of the model was affected more (or less) with regards to small-N design data of specific classes (components denoted as Ci). The results revealed that differences in accuracy were dependent on the component (Figure 2). Most classification errors occurred with C2 followed by C1, for both approaches. Moreover, though more variability was evident in A2, we observed that the predictive ability was overall higher with C3, C5 and C7 than with C1, C2 and C6.

As a result, starting from Chapter 4, we exclusively utilised datasets generated through Approach one for all subsequent experiments.

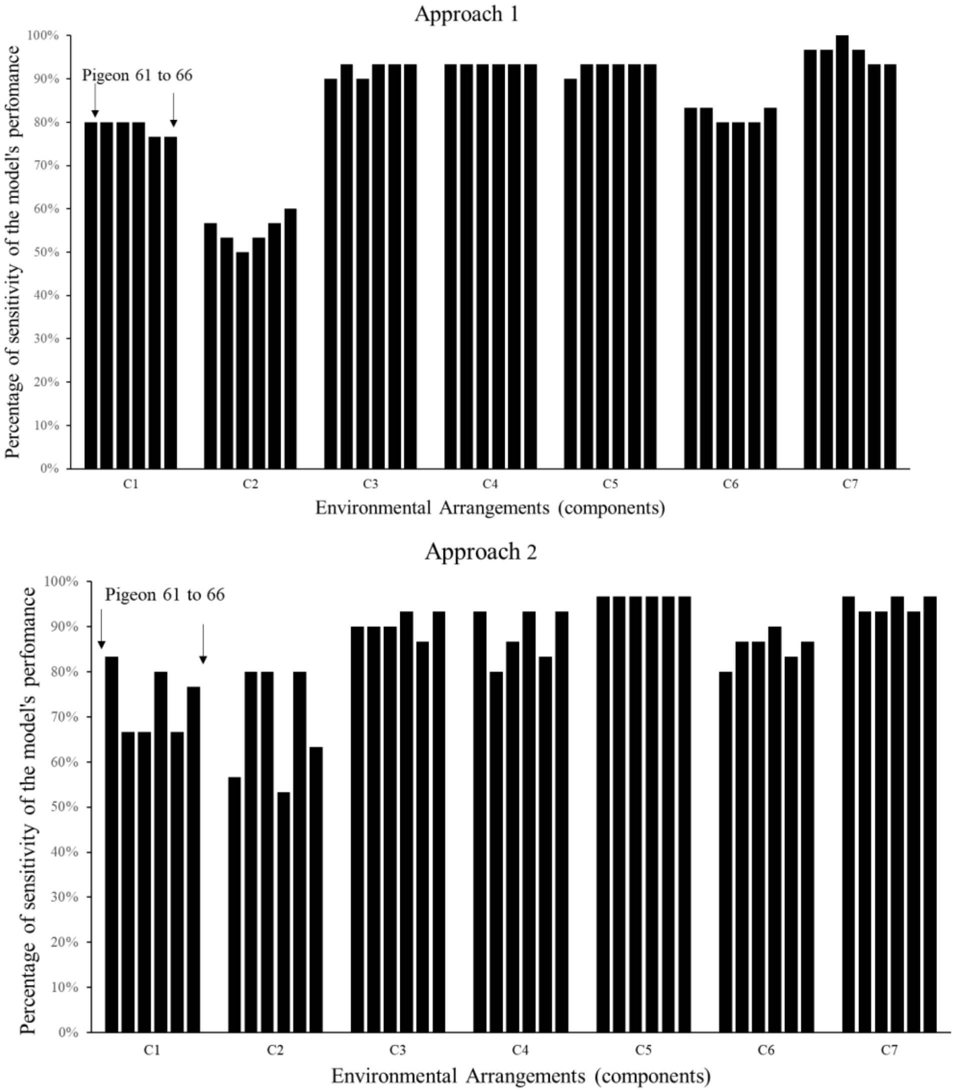

Figure 2. Percentage of classification rate (sensitivity) for all constructed models' predictions across all components. In each component, the presentation starts with the results when Pigeon 61 are out for testing through Pigeon 66.

B2. An example of the stepwise process extracting the data from the source study

Examples of the Steps Taken to Extract Pigeon Data from the Original Study (Landon & Davison, 2001) as generated by MED-PC⍰ software

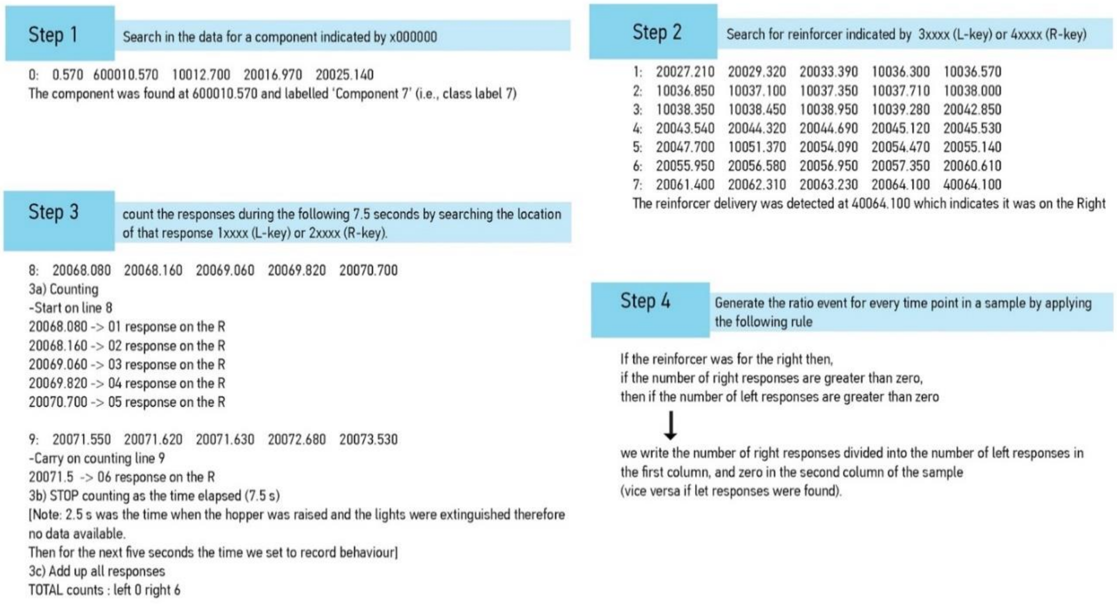

Step 1 Search in the data for a component indicated by x000000

```
0:   0.570  600010.570   10012.700   20016.970   20025.140
```

The component was found at 600010.570 and labelled 'Component 7' (i.e., class label 7)

Step 2 Search for reinforcer indicated by 3xxxx (L-key) or 4xxxx (R-key)

```
1:  20027.210  20029.320  20033.390  10036.300  10036.570
2:  10036.850  10037.100  10037.350  10037.710  10038.000
3:  10038.350  10038.450  10038.950  10039.280  20042.850
4:  20043.540  20044.320  20044.690  20045.120  20045.530
5:  20047.700  10051.370  20054.090  20054.470  20055.140
6:  20055.950  20056.580  20056.950  20057.350  20060.610
7:  20061.400  20062.310  20063.230  20064.100  40064.100
```

The reinforcer delivery was detected at 40064.100 which indicates it was on the Right

Step 3 count the responses during the following 7.5 seconds by searching the location of that response 1xxxx (L-key) or 2xxxx (R-key).

```
8:   20068.080   20068.160   20069.060   20069.820   20070.700
```

3a) Counting
-Start on line 8
20068.080 -> 01 response on the R
20068.160 -> 02 response on the R
20069.060 -> 03 response on the R
20069.820 -> 04 response on the R
20070.700 -> 05 response on the R

```
9:   20071.550   20071.620   20071.630   20072.680   20073.530
```

-Carry on counting line 9
20071.5 -> 06 response on the R
3b) STOP counting as the time elapsed (7.5 s)
[Note: 2.5 s was the time when the hopper was raised and the lights were extinguished therefore no data available.
Then for the next five seconds the time we set to record behaviour]
3c) Add up all responses
TOTAL counts : left 0 right 6

Step 4 Generate the ratio event for every time point in a sample by applying the following rule

If the reinforcer was for the right then,
if the number of right responses are greater than zero,
then if the number of left responses are greater than zero

↓

we write the number of right responses divided into the number of left responses in the first column, and zero in the second column of the sample
(vice versa if let responses were found).

Condition 1 and 6 events in the original experiment Landon and Davison (2001) were arranged on an IBM-PC compatible computer running MED-PC software. The computer recorded all events and the time they occurred within each Condition. We developed a computer code to identify all original experimental codes from the existing files and created four steps to extract pigeon data (Figure A1): (1) Searching in the data for a component indicated by x000000 that denote the start of each component (1-7), were x=0-6; (2) searching for each of the nine reinforcers as they were delivered in time by looking for the 3xxxx(L-key) or 4xxxx(R-key); (3) counting and adding the responses during the following 7.5 seconds (2.5 seconds for the hopper raised and 5 seconds of recording response frequency to the L or R key); and (4) generating the ratio event for every 5-second period. In this way, we extracted all data for input samples. The final outcome was a ratio that was a frequency event, which was created by calculating the ratio of left and right preferences based on the time window (5-sec) and on the actual location (L or R) where the reinforcer was delivered.

For example, let’s hypothesise that after taking steps 1-3 as presented in Figure A1, we extracted the following information: a pigeon was presented randomly with one of the seven different components (i.e., 27: 1, 9:1, 3:1, 1:1, 1:3, 1:9, and 1:27) and a reinforcer was allocated to the right, and after consuming the reinforcer, the pigeon then started pecking one time on the left-key and four times on the right-key during a period of five seconds.

The relative response of the pigeon would be 4/1 (the pigeon pecked four times more on the right key). In the reverse scenario were for example, the reinforcer was allocated on the left, and the pigeon pecked one time on the right and three times on the left, then the relative response would be 3/1 (the pigeon pecked three times on the left key). By using this method, all responses extracted from the original experiment into the samples occurred after the reinforcer was delivered for five seconds and were dependent on the reinforcer location in time.

B3. Single Spiking Neural Network Architecture for this Study

(This section was published as supplementary material for Plessas et al. (2022)).

Based on the structure of the perceptron, the single neuron model constructed to classify components based on pigeon choices, was formed by:

1. A Leaky Integrate-and-Fire (LIF) neuron as the temporal activation unit (Gerstner & Kistler, 2002c);
2. An m-by-2 matrix A of 2 input vectors $x_1, x_2 \in \mathbb{R}^m$ that represent the m pigeons' choices (either to the left x_1 or to the right x_2 key in different periods of 10, 7, 5 days;
3. A vector of real values $w \in \mathbb{R}^2 \mapsto [-0.1, 0.1]$ that represents the synaptic connections (‘weights’) between the input nodes and the LIF neuron

In Figure 1, a schematic representation of the classification process shows the input currents produced by samples of the same component having similar firing rates, whereas input currents of different components produce firing rates different enough to discriminate among the various components. Therefore, an input current $(\mathrm{I} = \mathrm{A} \cdot \mathrm{w})$, where *I* is the injected current, *A* the input vectors and *w* the neuron’s weight stimulated the neuron model to produce a sequence of artificial spikes with a specific firing rate. Samples were decomposed into two input vectors that represented pigeons’ left and right choices in periods of ten sessions (Figure 1).

Figure B1

The SNN Architecture Constructed for Classification of Seven Components Based on Pigeons' Choices

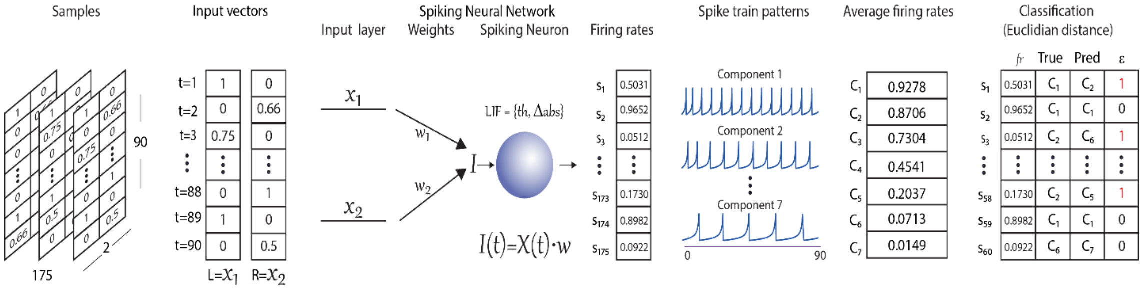

An input current $\mathrm{I} = \mathrm{A} \cdot \mathrm{w}$ stimulates the neuron model to produce a sequence of spikes with a specific firing rate. We trained the neuron to classify the components by adapting the weights w and generating the input current I that characterises each component. This approach is addressed as an optimisation process, where the objective is to find the optimal w values that maximise the classification accuracy. For this purpose, we applied the Differential Evolution (DE) algorithm (DE/rand/1/bin), a population-based heuristic method for global optimisation over continuous spaces (Storn & Price, 1997).

In this research, every candidate solution (also called an 'agent') of the DE population represents the weights w, for which the fitness function is the classification accuracy. To calculate the classification accuracy, we first compute the average firing rate of the samples belonging to each class, there being seven firing rates (seven components) in this particular case. Then, we label the retrodict component of each sample with the component having the closest firing rate. Finally, we compare the retrodicted and actual components of all samples and calculate the errors and then the classification accuracy.

In other words, every candidate solution of the population was a unique random vector $a = [a_1, a_2, a_3, a_4]$ mapped to the vector of weights $a_1 \mapsto w_1$, $a_2 \mapsto w_2$, and the LIF neuron's threshold $a_3 \mapsto th$ and refractory time $a_4 \mapsto rt$ (Figure 2).

Figure B2

The Differential Evolution Algorithm Architecture for Classification of Seven Components Based on Pigeons' Choices

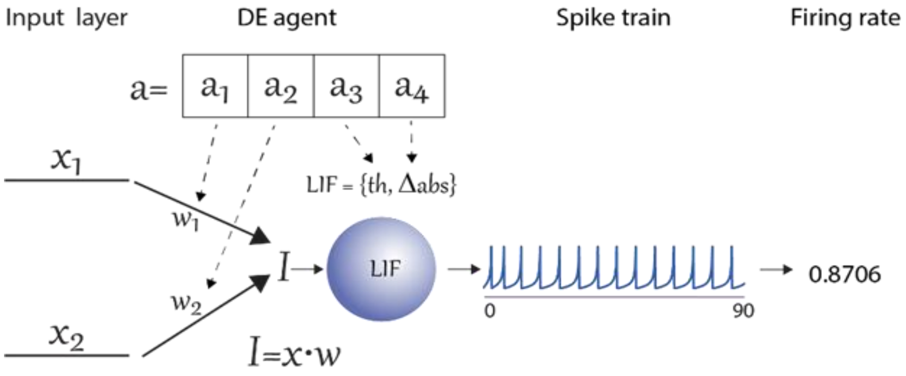

Training the Single Spiking Neural Network for this Study

Besides adapting the weights w, in this study we also optimised the LIF threshold and refractory time in applying the same DE algorithm. Both the weights and the LIF parameters formed an individual agent $x \in \mathbb{R}^4$. The weights were mapped to two elements x of each agent $w_1 \mapsto x_1, w_2 \mapsto x_2$ and the firing threshold and refractory period to its last two elements $th \mapsto x_3, rt \mapsto x_4$. Table 1 summarises the parameters used in this process.

Learning

Every training set sample s_i produced a vector of currents $I_i \in \mathbb{R}^n$ by the product $I_i = X_i \cdot w$, where X_i is the ith-sample of pigeon's choices and w the vector of weights (note that X is an n-by-2 matrix and w n-by-1 column vector). Every vector I_i was used as the injected current that stimulated the spiking neuron (with parameters th and Δ_{abs}) to produce a spike train $r_i \in \{0,1\}^n$ (vector of ones and zeros). We used the firing rate $f_i = \frac{1}{k_i}$ as a feature for classification, where k_i is the number of spikes (ones) of the spike train r_i.

Subsequently, we calculated the average firing rate of the training samples belonging to the same class and associated them with the respective component. Since every agent of the DE population was unique, it was impossible to know a priori the set of average firing rates associated to each of the components. Therefore, this part of the training process can be considered unsupervised learning. Like the unsupervised k-means clustering, this vector quantisation method aims to partition i observations (i samples) into k clusters (7 components) in which each observation belongs to the cluster with the nearest mean (cluster centers or cluster centroid), serving as a prototype of the cluster.

Since all the training samples are labelled, we can consider this final stage of the algorithm as supervised learning and calculate the training classification accuracy of the agent.

We can observe that every agent represents a unique SNN. We estimated the skill of every agent generated in all generations of the DE by means of a k-fold cross-validation applied to the training data.

Table B1

Parameters Utilised for Training and LIF's Parameters Optimisation

Method	Parameters	Limits
DE	Agent	4
	Population Size	40
	Max generations	150
	Crossover rate	0.7
	Mutation factor rate	0.1
SNN	Weight Range of values	[-0.01,0.01]
	LIF Threshold	[0.01, 1.0]
	LIF Absolute refractory time Δ_{abs}	[2,10]
K-fold	K (cross-validation)	5

Every agent represents a unique SNN. We estimated the skill We estimated the skill of every agent produced during the execution of the DE by means of a five-fold cross-validation applied to the data for training.

Figure 3 presents the algorithm validation procedure that ensured that every SNN model built produced similar accurate results with low standard deviation.

Figure B3

The Validation Procedure Used for the Algorithm

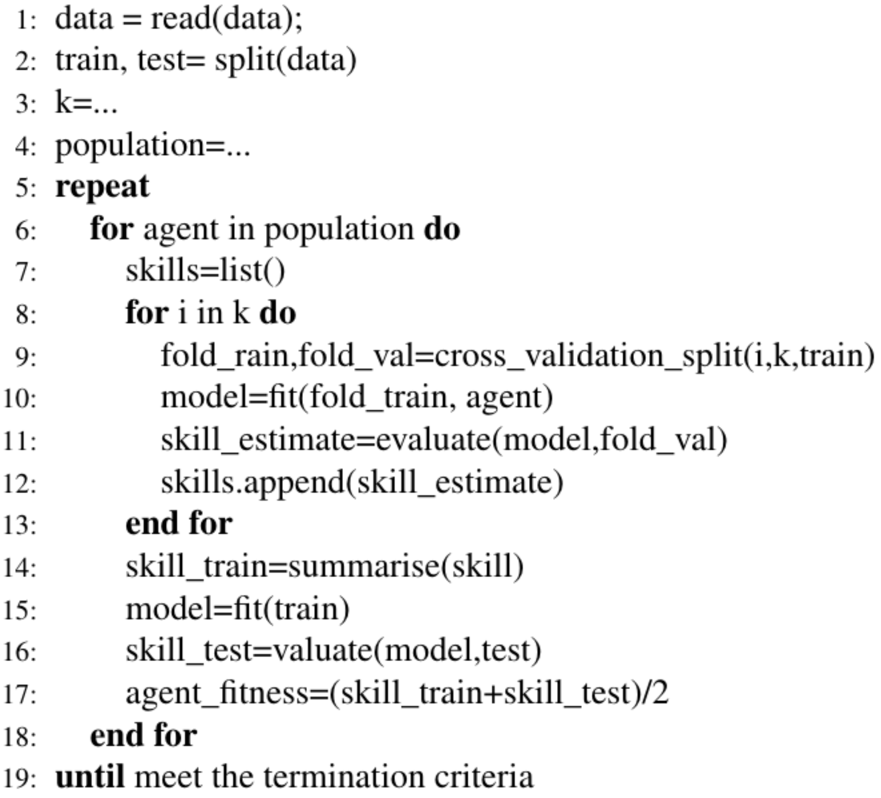

```
data = read(data);
train, test= split(data)
k=...
population=...
repeat
    for agent in population do
        skills=list()
        for i in k do
            fold_rain,fold_val=cross_validation_split(i,k,train)
            model=fit(fold_train, agent)
            skill_estimate=evaluate(model,fold_val)
            skills.append(skill_estimate)
        end for
        skill_train=summarise(skill)
        model=fit(train)
        skill_test=valuate(model,test)
        agent_fitness=(skill_train+skill_test)/2
    end for
until meet the termination criteria
```

B4. Pigeon Performances to SNN model performances (extended)

We investigated the pigeons' responses during the initial experimentation (Landon & Davison, 2001) based on the logarithmic version of the GML; the results were compared to the machine's outcomes. The data for analysis consisted of the choices pigeons made in the 5 seconds following the delivery of a reinforcer in each component. The data were aggregated according to the left-key over right-key response ratios; logarithms of these ratios were calculated and plotted as a function of each separate component given by the generalised matching law formula. Table 5 shows the values of sensitivity, bias, and the variance accounted for fits of the GML model for all 50 sessions. The analysis for Pigeon 61 and Pigeon 62 showed the highest sensitivity to reinforcement (0.93), followed by birds 63 and 64, which illustrates that their choice either to the left- or the right-key was mainly a function of the reinforcer ratio. Furthermore, the linear regression analysis for 5-sec data revealed that a consistent preference towards beyond that expected according to the GML the right alternative was evident across all pigeons in all components. This bias (log c) was relatively low. In contrast, with Pigeon 65 and 66 sensitivity was evidently lower. Pigeon 63 caused the highest bias followed by Pigeon 63.

Table 5
Sensitivity to reinforcement after the 9th successive component reinforcer for each pigeon. Also shown are respectively logc and r^2 values.

Participants	a	logc	r^2
Pigeon 61	0.936	0.023	0.997
Pigeon 62	0.937	-0.186	0.996
Pigeon 63	0.917	-0.611	0.982
Pigeon 64	0881	-0.144	0.987
Pigeon 65	0.731	-0.186	0.99
Pigeon 66	0.852	-0.254	0.997
Group	0.847	-0.188	0.998

Pigeons' choice behaviour was reflected in the SNN model outcomes. Overall, when sensitivity (a) to reinforcement was high and bias (log c) was low (Pigeon 1), high performance was found in the overall metrics. However, when a was lowest and log c remained high (Pigeon 6), we saw that the model's predictive ability declined.

To examine within-session variation effects of reinforcers, we analysed how preference changed within each component when data were aggregated according to the left-key over the right-key response ratio for only the first 5 sec after each delivery of reinforcer; the logarithms of these ratios were calculated and plotted as a function of successive reinforcers for each component seperately (Figure 4). In the original study it was reported that pigeons' preference to the left or right shifted along with changes in reinforcer ratios. Our analysis captured a snapshot of the original data by including the bird responses emitted for the first 5 seconds following each reinforcer delivery and it revealed the same pattern as the original study.

Figure 4 shows that choice behaviour shifted towards the alternative choice with the higher reinforcement. In particular, the biggest change was reported after the seventh reinforcer in C7, where preference was stabilised towards the left for all pigeons. Less discriminable reinforcer ratios within a component did not produce preferences as solid as that one. The smallest change was observed with C2 followed by C5.

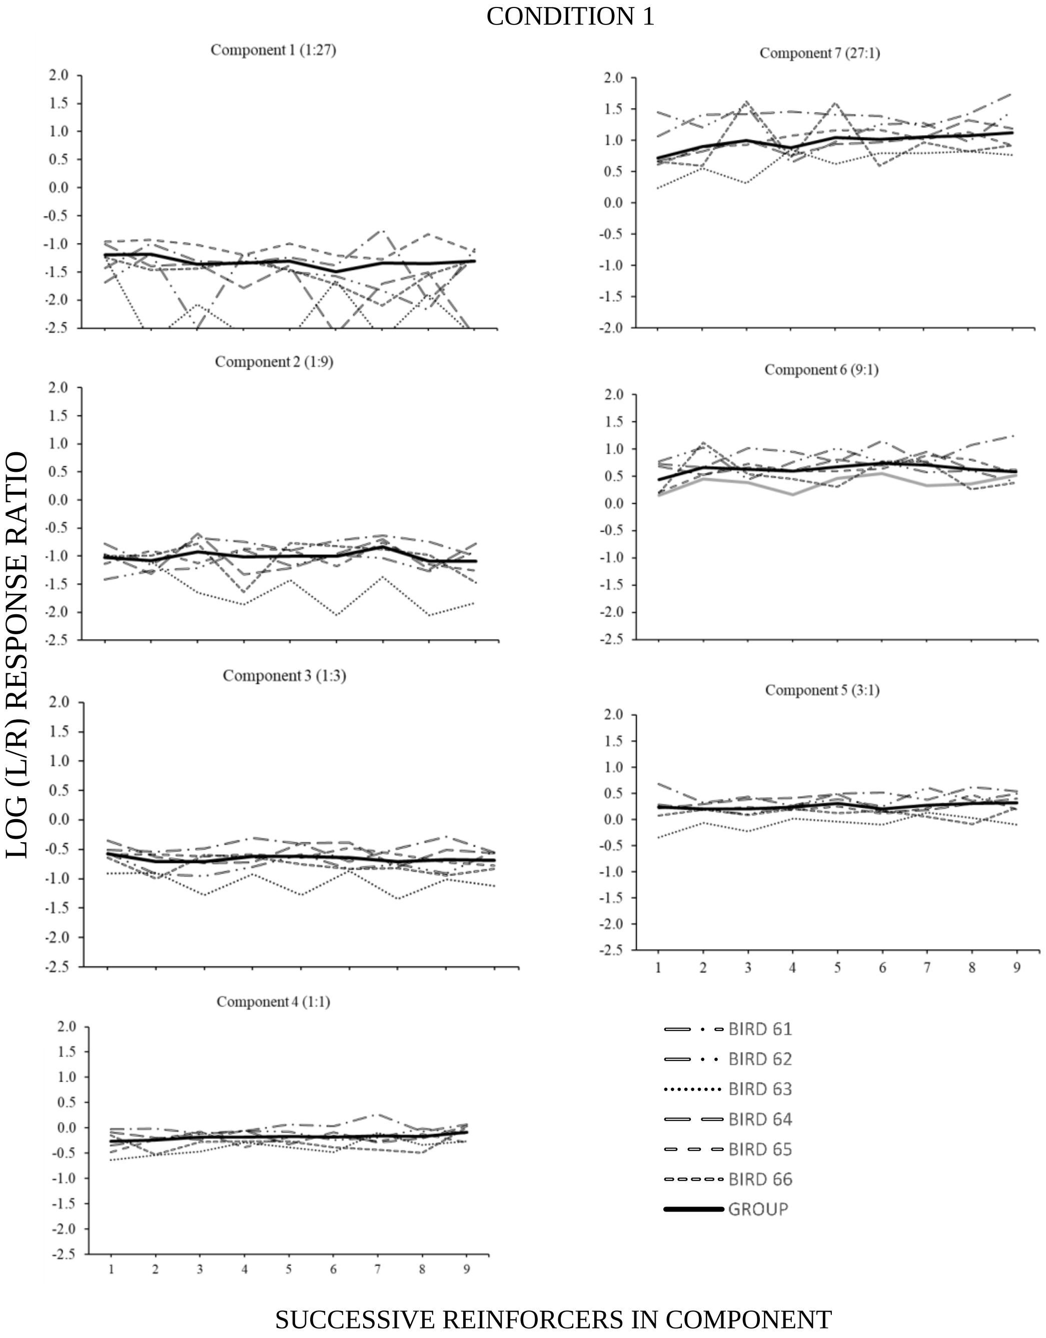

Figure 4. Log response ratios emitted for the first 5 seconds following each successive reinforcer delivery in each of the seven components of condition 1 of the original experiment. Response numbers for each alternative were summed across 50 sessions. On each graph, the separate plot show data from each individual subset and the group. On the x-axis, the nine successive reinforcers chosen for analysis are presented.

The analysis of classification per class label (see Results in Chapter 4) revelead that classification tasks were performed best with C7 and worst with C2. Pigeons' behaviour showed a change in preference in C7 and we can assume this move in choice was sufficient for the model to detect a pattern. On the other hand, in C2, the model could not identify or correlate preference to the right-key explicitly to reinforcement. It seems that if log c was zero, then possibly the machine could make a

100% accurate prediction. However, inherent bias interferes with the observable patterns. Possibly, the increase in sensitivity to reinforcement by itself may not ensure better predictions when influences from bias are present. Yet, the less sensitivity pigeons show to reinforcement, the harder it is for the SNN model to identify patterns and make predictions.

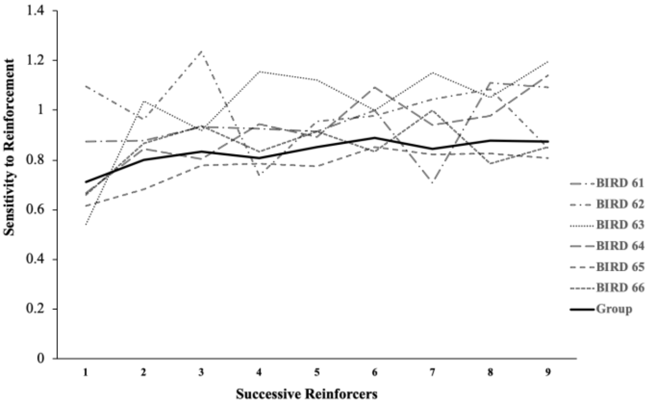

Figure 5. Sensitivity to reinforcement as a function of successive component reinforcers calculated using data from all 50 sessions.

As shown in Figure 5, we analysed sensitivity to reinforcement for each successive reinforcer. In the original study, sensitivity to reinforcers increased successively up to the 9th reinforcer. The effects of learning were significant and driven from the reinforcer ratios and the sequential effects of the reinforcer within the component. In our analysis, the sensitivity values started high, which can be explained with the small window of observation. Pigeons could be reusing the just-reinforced response rather than be controlled by the overall reinforcer ratio set in the component. Moreover, between-pigeon variability is high, indicating that each pigeon responded differently to reinforcement, which also reflects on the variability of the machine's metrics in the testing phase. The fact that the model detected a pattern and classified component-related choices indicates that the pigeons' responses are not just 'reinforcer pulses', i.e., repetitions of a just-reinforced response. The temporal proximity to the just-reinforced response should have a significant effect on choice, thus controlling subsequent responses. However, if they were just copying the just-reinforced response, there would be no inherent bias, their responses would strictly match reinforcement and components would be easy to discriminate.

Appendix C: Chapter 6 Supplementary Materials

C1. Baseline models for traditional approach

Following is the list of best single-neuron SNN developed to be used for traditional approach to generalisation tests.

Model 1: Model run 7
Model 2: Model run 10
Model 3: Model run 3
Model 4: Model run 8
Model 5: Model run 7
Model 6: Model run 8

Table C1. Average firing rates used for the traditional approach

Class (component)	P1-run7	P2-10	P3-3	P4-8	P5-7	P6-run8
1	0.90518519	0.90474	0.905925926	0.906123457	0.91930864	0.927851852
2	0.85107407	0.84742	0.84745679	0.848604938	0.85201235	0.870679012
3	0.72038272	0.71338	0.722407407	0.724234568	0.7252716	0.73045679
4	0.46191358	0.45204	0.462716049	0.449419753	0.44348148	0.454148148
5	0.22754321	0.21777	0.212592593	0.214259259	0.1994321	0.203777778
6	0.10159259	0.08385	0.090617284	0.089925926	0.06720988	0.071345679
7	0.03090123	0.02088	0.020234568	0.021407407	0.01408642	0.01491358

C2. Additional results for traditional and personalised modelling approach

Tables listed below present the models performances across traditional approach to generalisation and personalised modelling for the group of the participants.

Table C2

Generalisation results for personalised modelling and traditional approach dataset extracted from Condition 1 (Davison & Baum, 2000)

	Model 1		Model 2		Model 3		Model 4		Model 5		Model 6	
	P	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**
Accuracy	0.88	0.70	0.86	0.82	0.86	0.83	0.86	0.82	0.86	0.82	0.85	0.80
Recall	0.67	0.40	0.64	0.57	0.64	0.59	0.63	0.57	0.64	0.57	0.63	0.54
Specificity	0.92	0.80	0.92	0.89	0.91	0.90	0.91	0.89	0.91	0.89	0.91	0.87
Precision	0.67	0.40	0.64	0.57	0.64	0.59	0.63	0.57	0.64	0.57	0.63	0.54
F1	0.67	0.40	0.64	0.57	0.64	0.59	0.63	0.57	0.64	0.57	0.63	0.54
Informedness	0.60	0.21	0.56	0.46	0.55	0.49	0.54	0.45	0.55	0.45	0.53	0.41

Note. Personalised modelling results are denoted as ‘P’ and traditional approach as ‘Tr’

Table C3

Generalisation results for personalised modelling and traditional approach extracted from Condition 4 (Davison & Baum, 2000)

	Model 1		Model 2		Model 3		Model 4		Model 5		Model 6	
	P	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**
Accuracy	0.91	0.90	0.90	0.90	0.86	0.89	0.91	0.89	0.91	0.92	0.90	0.90
Recall	0.74	0.71	0.72	0.71	0.64	0.69	0.73	0.69	0.74	0.76	0.73	0.72
Specificity	0.95	0.94	0.94	0.94	0.91	0.93	0.94	0.93	0.94	0.95	0.94	0.94
Precision	0.74	0.71	0.72	0.71	0.55	0.69	0.73	0.69	0.74	0.76	0.73	0.72
F1	0.74	0.71	0.72	0.71	0.64	0.69	0.73	0.69	0.74	0.76	0.73	0.72
Informedness	0.69	0.65	0.66	0.65	0.64	0.62	0.68	0.62	0.68	0.71	0.67	0.66

Note. Personalised modelling results are denoted as ‘P’ and traditional approach as ‘Tr’

Table C4

Generalisation results for personalised modelling and traditional approach from Condition 9 (Davison & Baum, 2000)

	Model 1		Model 2		Model 3		Model 4		Model 5		Model 6	
	P	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**
Accuracy	0.92	0.90	0.91	0.89	0.92	0.90	0.90	0.89	0.90	0.88	0.88	0.88
Recall	0.76	0.73	0.73	0.71	0.76	0.73	0.73	0.69	0.73	0.68	0.68	0.68
Specificity	0.95	0.94	0.94	0.94	0.95	0.94	0.94	0.93	0.94	0.93	0.93	0.93
Precision	0.76	0.73	0.73	0.71	0.76	0.73	0.73	0.69	0.73	0.68	0.68	0.68
F1	0.76	0.73	0.73	0.71	0.76	0.73	0.73	0.69	0.73	0.68	0.68	0.68
Informedness	0.70	0.67	0.67	0.64	0.70	0.67	0.67	0.62	0.67	0.61	0.61	0.61

Note. Personalised modelling results are denoted as ‘P’ and traditional approach as ‘Tr’

C3. Preliminary results using human and rat datasets

The approach underwent testing using datasets from two behavioural paradigms that used other mammals as participants, humans and rat. The two studies employed a consistent training methodology for teaching the humans and rats to make binary choices. The datasets were extracted from experimental conditions as the one outlined in Plessas et al. (2022). Furthermore, their choice behaviour was extracted by the same experimental conditions utilised during the development of the single-neuron SNN model.

Table C5

Samples Extracted by All Mammals Across Studies

Species	Samples		Points of 5-sec periods	Total
Pigeons (n=6) (D&B,2000)	24	(7 components*4 samples per pigeon)	81-90	168
Rats (n=8)	70	(7 components*10 samples per pigeon)	51-90	518
Human (n=6)	24	(7 components*1 samples per pigeon)	48-60	42
Human* (n=6)	24	(7 components*1 samples per pigeon)	49	42

Note: Human* dataset was the same dataset as Human but included all choice behaviours made by the humans from the delivery of one reinforcer to the next, and not a small window (5-s) of data after the delivery of the reinforcer. This was done as humans produced smaller and variable datasets and wanted to assess performance in both conditions (all binary choices versus a snapshot of data).

Table C6

Generalisation results for personalised modelling and traditional approach rat datasets extracted from Landon et al. (2007)

	Model 1		Model 2		Model 3		Model 4		Model 5		Model 6	
	P	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**
Accuracy	0.86	0.73	0.85	0.74	0.85	0.76	0.85	0.74	0.85	0.76	0.85	0.74
Recall	0.64	0.44	0.62	0.45	0.63	0.47	0.62	0.44	0.62	0.47	0.62	0.45
Specificity	0.91	0.83	0.91	0.83	0.91	0.84	0.91	0.83	0.91	0.84	0.91	0.83
Precision	0.64	0.44	0.62	0.45	0.63	0.47	0.62	0.44	0.62	0.47	0.62	0.45
F1	0.64	0.44	0.62	0.45	0.63	0.47	0.62	0.44	0.62	0.47	0.62	0.45
Informedness	0.55	0.27	0.53	0.28	0.54	0.31	0.53	0.27	0.53	0.32	0.52	0.27

Note. Personalised modelling results are denoted as ‘P’ and traditional approach as ‘Tr’

Table C7

Generalisation results for personalised modelling and traditional approach a snapshot of human datasets extracted from Krägeloh et a. (2010)

	Model 1		Model 2		Model 3		Model 4		Model 5		Model 6	
	P	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**
Accuracy	0.74	0.37	0.74	0.37	0.74	0.37	0.74	0.37	0.74	0.37	0.74	0.37
Recall	0.45	0.14	0.45	0.14	0.45	0.14	0.45	0.14	0.45	0.14	0.45	0.14
Specificity	0.83	0.50	0.83	0.50	0.83	0.50	0.83	0.50	0.83	0.50	0.83	0.50
Precision	0.45	0.14	0.45	0.14	0.45	0.14	0.45	0.14	0.45	0.14	0.45	0.14
F1	0.45	0.14	0.45	0.14	0.45	0.14	0.45	0.14	0.45	0.14	0.45	0.14
Informedness	0.28	-	0.28	-	0.28	-	0.28	-	0.28	-	0.28	-

Note. Personalised modelling results are denoted as ‘P’ and traditional approach as ‘Tr’. datasets The snapshot of data is for 5-sec periods after the delivery of each reinforcer.

Table C8

Generalisation results for personalised modelling and traditional approach human datasets extracted from Krägeloh *et al. (2010)*

	Model 1		Model 2		Model 3		Model 4		Model 5		Model 6	
	P	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**	**P**	**Tr**
Accuracy	0.81	0.58	0.78	0.61	0.79	0.64	0.78	0.61	0.81	0.70	0.81	0.68
Recall	0.55	0.29	0.50	0.31	0.52	0.33	0.50	0.31	0.55	0.40	0.55	0.38
Specificity	0.88	0.71	0.86	0.73	0.87	0.75	0.86	0.73	0.88	0.80	0.88	0.79
Precision	0.55	0.29	0.50	0.31	0.52	0.33	0.50	0.31	0.55	0.40	0.55	0.38
F1	0.55	0.29	0.50	0.31	0.52	0.33	0.50	0.31	0.55	0.40	0.55	0.38
Informedness	0.43	-	0.36	0.04	0.39	0.08	0.36	0.04	0.43	0.21	0.43	0.17

Note. Personalised modeling results are denoted as 'P' and traditional approach as 'Tr'. Human datasets include all data generated in the experiment.

Figure C1

F-1 scores on Performance for all Personalised Models with all group datasets

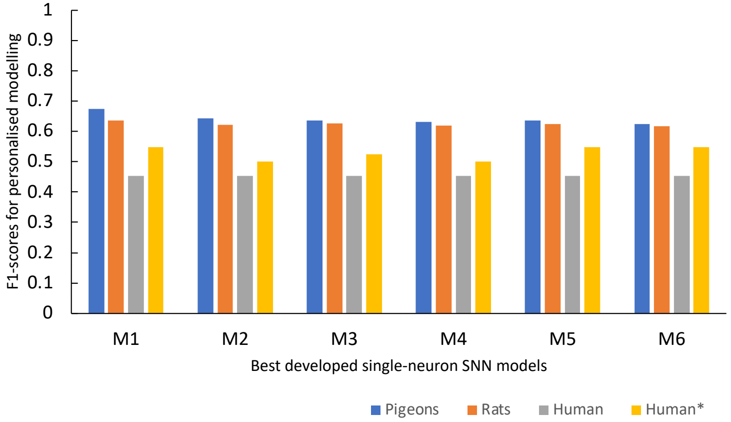

Note: Datasets extracted from Condition 1 in Landon & Davison (2001), Landon et al. (2007) and Krägeloh et a. (2010)
Human* dataset was the same dataset as Human but included all choice behaviours made by the humans from the delivery of one reinforcer to the next, and not a small window (5-s) of data after the delivery of the reinforcer (denoted as Human).

C4. Preliminary results for individual datasets used for personalised modelling

Tables listed below present some preliminary experimentation to investigate the personalised modelling approach ability to make retrodictions with individual datasets.

Table C9

Presenting an example on how all models performed for each individual pigeon dataset extracted from Condition 1 (Davison & Baum, 2000)

	Model-1		**Model-2**		**Model-3**		**Model-4**		**Model-5**		**Model-6**	
Data set Pigeon 91	Accuracy	0.91	Accuracy	0.93	Accuracy	0.93	Accuracy	0.93	Accuracy	0.90	Accuracy	0.90
	Recall	0.75	Recall	0.79	Recall	0.79	Recall	0.79	Recall	0.71	Recall	0.71
	Specificity	0.95	Specificity	0.96	Specificity	0.96	Specificity	0.96	Specificity	0.94	Specificity	0.94
	Precision	0.75	Precision	0.79	Precision	0.79	Precision	0.79	Precision	0.71	Precision	0.71
	F1	0.75	F1	0.79	F1	0.79	F1	0.79	F1	0.71	F1	0.71
	Informedness	0.70	Informedness	0.74	Informedness	0.74	Informedness	0.74	Informedness	0.65	Informedness	0.65
Data set Pigeon 92	Accuracy	0.95	Accuracy	0.94	Accuracy	0.94	Accuracy	0.94	Accuracy	0.91	Accuracy	0.91
	Recall	0.86	Recall	0.82	Recall	0.82	Recall	0.82	Recall	0.75	Recall	0.75
	Specificity	0.97	Specificity	0.97	Specificity	0.97	Specificity	0.97	Specificity	0.95	Specificity	0.95
	Precision	0.86	Precision	0.82	Precision	0.82	Precision	0.82	Precision	0.75	Precision	0.75
	F1	0.86	F1	0.82	F1	0.82	F1	0.82	F1	0.75	F1	0.75
	Informedness	0.83	Informedness	0.79	Informedness	0.79	Informedness	0.79	Informedness	0.70	Informedness	0.70
Data set Pigeon 93	Accuracy	0.95	Accuracy	0.95	Accuracy	0.95	Accuracy	0.94	Accuracy	0.95	Accuracy	0.94
	Recall	0.86	Recall	0.86	Recall	0.86	Recall	0.82	Recall	0.86	Recall	0.82
	Specificity	0.97	Specificity	0.97	Specificity	0.97	Specificity	0.97	Specificity	0.97	Specificity	0.97
	Precision	0.86	Precision	0.86	Precision	0.86	Precision	0.82	Precision	0.86	Precision	0.82
	F1	0.86	F1	0.86	F1	0.86	F1	0.82	F1	0.86	F1	0.82
	Informedness	0.83	Informedness	0.83	Informedness	0.83	Informedness	0.79	Informedness	0.83	Informedness	0.79

	Model-1		Model-2		Model-3		Model-4		Model-5		Model-6	
Data set Pigeon 94	Accuracy	0.93	Accuracy	0.91	Accuracy	0.91	Accuracy	0.91	Accuracy	0.93	Accuracy	0.93
	Recall	0.79	Recall	0.75	Recall	0.75	Recall	0.75	Recall	0.79	Recall	0.79
	Specificity	0.96	Specificity	0.95	Specificity	0.95	Specificity	0.95	Specificity	0.96	Specificity	0.96
	Precision	0.79	Precision	0.75	Precision	0.75	Precision	0.75	Precision	0.79	Precision	0.79
	F1	0.79	F1	0.75	F1	0.75	F1	0.75	F1	0.79	F1	0.79
	Informedness	0.74	Informedness	0.70	Informedness	0.70	Informedness	0.70	Informedness	0.74	Informedness	0.74
Data set Pigeon 95	Accuracy	0.95	Accuracy	0.94	Accuracy	0.94	Accuracy	0.94	Accuracy	0.97	Accuracy	0.97
	Recall	0.86	Recall	0.82	Recall	0.82	Recall	0.82	Recall	0.89	Recall	0.89
	Specificity	0.97	Specificity	0.97	Specificity	0.97	Specificity	0.97	Specificity	0.98	Specificity	0.98
	Precision	0.86	Precision	0.82	Precision	0.82	Precision	0.82	Precision	0.89	Precision	0.89
	F1	0.86	F1	0.82	F1	0.82	F1	0.82	F1	0.89	F1	0.89
	Informedness	0.83	Informedness	0.79	Informedness	0.79	Informedness	0.79	Informedness	0.87	Informedness	0.87
Data set Pigeon 96	Accuracy	0.93	Accuracy	0.91	Accuracy	0.86	Accuracy	0.91	Accuracy	0.88	Accuracy	0.88
	Recall	0.79	Recall	0.75	Recall	0.64	Recall	0.75	Recall	0.68	Recall	0.68
	Specificity	0.96	Specificity	0.95	Specificity	0.92	Specificity	0.95	Specificity	0.93	Specificity	0.93
	Precision	0.79	Precision	0.75	Precision	0.64	Precision	0.75	Precision	0.68	Precision	0.68
	F1	0.79	F1	0.75	F1	0.64	F1	0.75	F1	0.68	F1	0.68
	Informedness	0.74	Informedness	0.70	Informedness	0.56	Informedness	0.70	Informedness	0.61	Informedness	0.61

Table C10. Presenting an example on how all models performed for each individual rat dataset extracted from Condition 1 (Landon et al., 2007)

	Model-1		**Model-2**		**Model-3**		**Model-4**		**Model-5**		**Model-6**	
Data set Rat 1	Accuracy	0.88	Accuracy	0.88	Accuracy	0.87	Accuracy	0.89	Accuracy	0.88	Accuracy	0.88
	Recall	0.69	Recall	0.67	Recall	0.66	Recall	0.70	Recall	0.67	Recall	0.67
	Specificity	0.93	Specificity	0.92	Specificity	0.92	Specificity	0.93	Specificity	0.92	Specificity	0.92
	Precision	0.69	Precision	0.67	Precision	0.66	Precision	0.70	Precision	0.67	Precision	0.67
	F1	0.69	F1	0.67	F1	0.66	F1	0.70	F1	0.67	F1	0.67
	Informedness	0.61	Informedness	0.60	Informedness	0.58	Informedness	0.63	Informedness	0.60	Informedness	0.60
Data set Rat 2	Accuracy	0.88	Accuracy	0.88	Accuracy	0.88	Accuracy	0.88	Accuracy	0.88	Accuracy	0.88
	Recall	0.69	Recall	0.67	Recall	0.69	Recall	0.69	Recall	0.69	Recall	0.67
	Specificity	0.93	Specificity	0.92	Specificity	0.93	Specificity	0.93	Specificity	0.93	Specificity	0.92
	Precision	0.69	Precision	0.67	Precision	0.69	Precision	0.69	Precision	0.69	Precision	0.67
	F1	0.69	F1	0.67	F1	0.69	F1	0.69	F1	0.69	F1	0.67
	Informedness	0.61	Informedness	0.60	Informedness	0.61	Informedness	0.61	Informedness	0.61	Informedness	0.60
Data set Rat 3	Accuracy	0.90	Accuracy	0.87	Accuracy	0.87	Accuracy	0.86	Accuracy	0.85	Accuracy	0.85
	Recall	0.73	Recall	0.66	Recall	0.66	Recall	0.63	Recall	0.61	Recall	0.61
	Specificity	0.94	Specificity	0.92	Specificity	0.92	Specificity	0.91	Specificity	0.91	Specificity	0.91
	Precision	0.73	Precision	0.66	Precision	0.66	Precision	0.63	Precision	0.61	Precision	0.61
	F1	0.73	F1	0.66	F1	0.66	F1	0.63	F1	0.61	F1	0.61
	Informedness	0.67	Informedness	0.58	Informedness	0.58	Informedness	0.54	Informedness	0.52	Informedness	0.52
Data set Rat 4	Accuracy	0.88	Accuracy	0.88	Accuracy	0.90	Accuracy	0.88	Accuracy	0.89	Accuracy	0.89
	Recall	0.69	Recall	0.67	Recall	0.71	Recall	0.67	Recall	0.70	Recall	0.70
	Specificity	0.93	Specificity	0.92	Specificity	0.94	Specificity	0.92	Specificity	0.93	Specificity	0.93
	Precision	0.69	Precision	0.67	Precision	0.71	Precision	0.67	Precision	0.70	Precision	0.70
	F1	0.69	F1	0.67	F1	0.71	F1	0.67	F1	0.70	F1	0.70
	Informedness	0.61	Informedness	0.60	Informedness	0.65	Informedness	0.60	Informedness	0.63	Informedness	0.63

	Model-1		Model-2		Model-3		Model-4		Model-5		Model-6	
Data set Rat 5	Accuracy	0.83	Accuracy	0.81	Accuracy	0.82	Accuracy	0.82	Accuracy	0.83	Accuracy	0.83
	Recall	0.59	Recall	0.56	Recall	0.57	Recall	0.57	Recall	0.59	Recall	0.59
	Specificity	0.89	Specificity	0.88	Specificity	0.89	Specificity	0.89	Specificity	0.89	Specificity	0.89
	Precision	0.59	Precision	0.56	Precision	0.57	Precision	0.57	Precision	0.59	Precision	0.59
	F1	0.59	F1	0.56	F1	0.57	F1	0.57	F1	0.59	F1	0.59
	Informedness	0.48	Informedness	0.44	Informedness	0.46	Informedness	0.46	Informedness	0.48	Informedness	0.48
Data set Rat 6	Accuracy	0.90	Accuracy	0.89	Accuracy	0.90	Accuracy	0.88	Accuracy	0.86	Accuracy	0.86
	Recall	0.71	Recall	0.70	Recall	0.71	Recall	0.69	Recall	0.63	Recall	0.63
	Specificity	0.94	Specificity	0.93	Specificity	0.94	Specificity	0.93	Specificity	0.91	Specificity	0.91
	Precision	0.71	Precision	0.70	Precision	0.71	Precision	0.69	Precision	0.63	Precision	0.63
	F1	0.71	F1	0.70	F1	0.71	F1	0.69	F1	0.63	F1	0.63
	Informedness	0.65	Informedness	0.63	Informedness	0.65	Informedness	0.61	Informedness	0.54	Informedness	0.54
Data set Rat 7	Accuracy	0.92	Accuracy	0.91	Accuracy	0.92	Accuracy	0.92	Accuracy	0.92	Accuracy	0.92
	Recall	0.77	Recall	0.74	Recall	0.77	Recall	0.76	Recall	0.76	Recall	0.76
	Specificity	0.95	Specificity	0.95	Specificity	0.95	Specificity	0.95	Specificity	0.95	Specificity	0.95
	Precision	0.77	Precision	0.74	Precision	0.77	Precision	0.76	Precision	0.76	Precision	0.76
	F1	0.77	F1	0.74	F1	0.77	F1	0.76	F1	0.76	F1	0.76
	Informedness	0.72	Informedness	0.69	Informedness	0.72	Informedness	0.71	Informedness	0.71	Informedness	0.71
Data set Rat 8	Accuracy	0.93	Accuracy	0.93	Accuracy	0.94	Accuracy	0.93	Accuracy	0.94	Accuracy	0.93
	Recall	0.79	Recall	0.79	Recall	0.82	Recall	0.79	Recall	0.82	Recall	0.79
	Specificity	0.96	Specificity	0.96	Specificity	0.97	Specificity	0.96	Specificity	0.97	Specificity	0.96
	Precision	0.79	Precision	0.79	Precision	0.82	Precision	0.79	Precision	0.82	Precision	0.79
	F1	0.79	F1	0.79	F1	0.82	F1	0.79	F1	0.82	F1	0.79
	Informedness	0.74	Informedness	0.74	Informedness	0.79	Informedness	0.74	Informedness	0.79	Informedness	0.74

Table C11. Presenting an example on how the model performed for the dataset of each human extracted from Condition 1 in Krägeloh et a. (2010)

	Human 1	**Human 2**	**Human 3**	**Human 4**	**Human 5**	**Human 6**
Accuracy	0.95	0.90	1.00	0.95	0.90	0.95
Recall	0.86	0.71	1.00	0.86	0.71	0.86
Specificity	0.97	0.94	1.00	0.97	0.94	0.97
Precision	0.86	0.71	1.00	0.86	0.71	0.86
F1	0.86	0.71	1.00	0.86	0.71	0.86
Informedness	0.83	0.65	1.00	0.83	0.65	0.83

Note: Human datasets when 5-sec of choice data were encountered after the delivery of the reinforcer

C5. Single spiking neural network architecture for this study

(This section was published as supplementary material for Plessas et al. (2022).

SNN architecture

We built an SNN based on the architecture described in (Vazquez & Cachón, 2010) and extended it to recognise temporal data (i.e., pigeon's binary choices). The architecture consisted of five main elements:

1. A computational unit based on the Leaky Integrate-and-Fire (LIF) neuron model (Gerstner & Kistler, 2002).

   $$\tau_m \frac{dv}{dt} = -v(t) + RI(t) \quad (1)$$

   *

   Where $v(t)$ represents the membrane potential at time t, τ_m is the membrane time constant, R is the resistance, and $I(t)$ is the input current at time t. When the membrane potential $v(t)$ reaches a threshold value th, it emits a spike and resets to a reset potential v_r. The neuron stays inactive during an absolute refractory period Δ_{abs}. Then, the leaky integration process is re-initiated following a delay of Δ_{abs} after the spike.
2. A set of input patterns $p = \{X^i, k\}_{i=1}$ where X^i is an n-by-2 matrix and each column vector $x_1^i, x_2^i \in \mathbb{R}^n$ contains the n pigeon's choices (either to the left x_1^i or to the right x_2^i key), and $k = 1,2, \dots, K$ is the class (class=reward arrangement or briefly component) to which X^i belongs.
3. A vector of weights $w \in \mathbb{R}^2$ that represents the "synaptic connections" between the input nodes and the neuron.
4. An input signal $I_i = X^i \cdot w$ that simulates an injected current that feds the neuron to produce a spike train $s \in \{0,1\}^n$ where 1 represents a spike (action potential) and 0 is no spike.
5. Lastly, for classification, we first calculated the firing rate of each input pattern as,

$$fr_i = \frac{spk}{|s_i|} \quad (2)$$

being spk the number of spikes. Then we calculated the average firing rate $F \in \mathbb{R}^k$ of all the samples belonging to each class k. Then, we determined the class of an input pattern X^i as the label of the class with the closest average firing rate using Equation (3)

$$c_i = arg\ min_{k=1}^{K}(|F_k - \ fr_i|) \quad (3)$$

SNN Training

We addressed the training as an optimisation process, where the objective function was to find the vector of weights w that maximises the classification accuracy of the input patterns p. In this study, we applied the Differential Evolution (DE) algorithm (DE/rand/1/bin), a population-based heuristic method for global optimisation over continuous spaces (Storn & Price, 1997).

To exert maximum leverage of the algorithm, we also included the LIF neuron's threshold th and absolute refractory time Δ_{abs} as elements of the candidate solution the DE population q. These two properties of the neuron affect the output spike train and the firing rate used for the classification. In this study, we define a population of candidate solutions (agents) as $q = \{a_1, a_2, \dots, a_{NP}\} \ \forall \ a_i = \{w_1, w_2, th, \Delta_{abs}\}$ where $NP = 10d$ is the population size and d the size of the agents (dimensionality of the problem). A representation of the single neuron architecture and the DE agent is shown in Figure 1. The models created using the proposed architecture and training strategy are referred to as baseline models.

Figure C2

A schematic representation of the product of an input and weights that produce an injected current that feds the spiking neuron to produce a spike train with a specific firing rate.

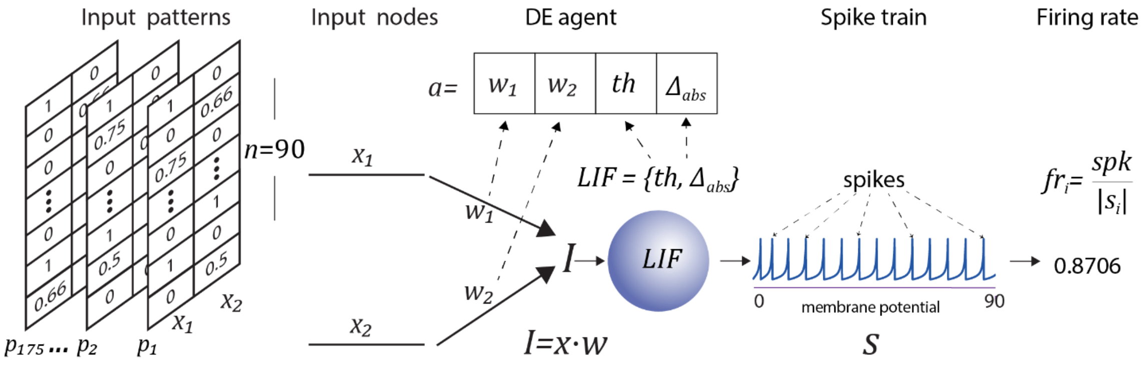

Traditional Approach to Generalisation

Let the $u_{base} = \{w \in \mathbb{R}^2, th, \Delta_{abs}\}$ be the candidate solution that produced the highest perfomance to be the underlying patterns of a pre-trained (baseline) model, p' the dataset of input patterns of a new source domain and $F \in \mathbb{R}^k$ the outputs (average firing rates) of a pre-trained (baseline) SNN model.

Then we use the calculated average firing rate F∈ R^k of all the samples belonging to each class k, generated by the pre-trained (baseline) model to determined the class of an input pattern X^i (using a the new set p'_{gen}) as the label of the class with the closest average firing rate using Equation (3).

3. $c_i = arg\ min_{k=1}^{K}(|F_k - fr_i|)$ (3)

The variable c_i represents the predicted class for the new set p'_{gen} The formula calulates the distance between the firing rate fr_i of the new set p'_{gen} and uses the firing rate of each K nearest neighbouts already calculated in baseline set (F_k). The distance metric used is the Euclidean distance, which is calculated as the square root of the sum of the squared differences between the corresponding features. The formula then selects the class that is most closest (frequent) among the K nearest neighbors as the predicted class c_i.

Personalised Modelling for SNN

Initially, we allowed again the $u_{base} = \{w \in \mathbb{R}^2, th, \Delta_{abs}\}$ to be the underlying patterns of a pre-trained (baseline) model, p' the dataset of input patterns of a new source domain and F' ∈ $\mathbb{R}^k$ the outputs (average firing rates) of a new SNN model. First, we transferred u_{base} to a new SNN. Then, the outputs F' of the SNN were adjusted by re-calculating the firing rates when using a the new set p'_{test}. The performance of the model was assessed using Equation (3) to classify the testing dataset p'_{test}, with F'_k. now

4. $c_i = arg\ min_{k=1}^{K}(|F'_k - fr_i|) \quad (4)$

This approach can be considered as a form of transfer learning, specifically feature extraction by leveraging the learnt representations from the pre-trained model (baseline) to improve the performance of the new SNN model on the new set of input patterns. However in this study we view this process as a simple adjustment because the underlying structure and weights of the SNN model are not being modified. Instead, only the firing rates are being adjusted to better fit the new input patterns.

Figure 2 shows a comparative diagram of how generalisation is tested using a traditional approach and personalised approach.

Figure C3

The weights and LIF parameters (in green) of the baseline model were used to for personlised SSN modelling on a new dataset.

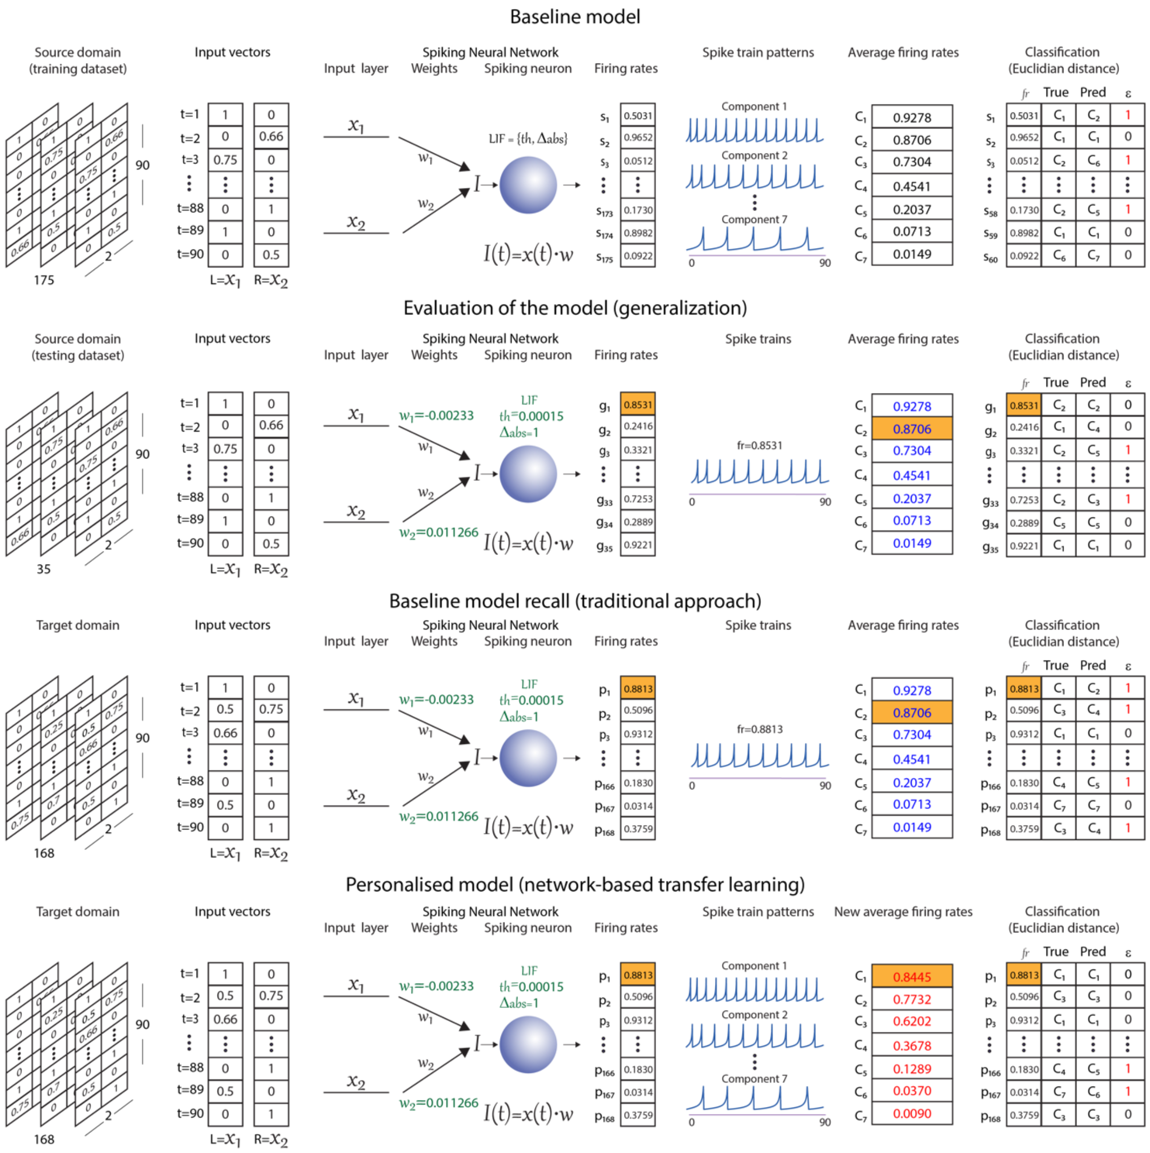

Note. Input patterns from different domains produced different average firing rates (in red), which were more suitable to the problem.