

IoTMindCare: Safe and Personalised IoT-Based Depression Management

Sanaz Zamani

Supervisors: Prof. Minh Nguyen
Prof. Samaneh Madanian
Prof. Roopak Sinha

School of Engineering, Computer & Mathematical Sciences
Auckland University of Technology

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

October 2026

Abstract

Depression is a complex and highly individual mental health condition that requires timely, personalised, and context-aware support. Traditional depression management often relies on periodic clinical assessments and self-reported symptoms, which can limit the ability to detect changes over time, respond to individual needs, and provide continuous support in everyday environments. Recent advances in Internet of Things (IoT)-enabled health technologies have created new opportunities for continuous and personalised depression management. However, existing IoT-based mental health systems remain fragmented, with many focusing on isolated sensing, symptom prediction, or prototype-level functionality rather than integrated, safety-aware, personalised, and implementation-informed solutions. This thesis proposes *IoTMindCare*, an implementation-informed reference architecture for safe, personalised, and context-aware depression monitoring and management. The thesis addresses these gaps by integrating evidence from existing IoT-based depression systems with data-driven monitoring, engagement-focused design, architectural modelling, and expert-informed refinement to inform safer, more personalised, and more implementation-aware depression management. First, a systematic review establishes the current state of IoT-based solutions for depression and emotional well-being, identifying key gaps in sensing approaches, data analysis, user engagement, privacy, security, interoperability, and system integration. Second, a data-driven IoT approach examines how environmental factors, particularly temperature and humidity, may provide an additional contextual layer for interpreting mental health-related patterns. Third, a gamified IoT-based study explores how micro-games, adaptive surveys, sensor data, and personalised feedback can support user engagement, self-awareness, and ongoing self-management. Fourth, the *IoTMindCare* reference architecture integrates these insights into a modular framework that combines multimodal sensing, hybrid edge-cloud processing, personalised analytics, anomaly detection, alert mechanisms, and stakeholder-centred design. Finally, an expert-informed refinement study extends *IoTMindCare* by incorporating technical and clinical perspectives on feasibility, safety, personalisation, governance, interoperability, workflow integration, and deployment constraints. The main contribution of this thesis is a coherent architectural foundation for moving IoT-based depression management beyond fragmented prototypes toward safer and more personalised digital mental health

systems. The architecture is evaluated through design-level analysis and expert-informed refinement; its integrated operational performance, clinical effectiveness, and deployment feasibility remain to be established through future empirical studies. Rather than treating depression monitoring as a purely predictive or technology-driven task, the thesis argues that future IoT-based depression management requires contextual interpretation, personalised calibration, human-in-the-loop oversight, alert governance, privacy-aware design, and alignment with clinical and real-world implementation requirements. *IoTMindCare* provides a foundation for future prototyping, simulation, pilot studies, and empirical validation of IoT-enabled systems for depression monitoring and management.

Table of contents

List of figures **xi**

List of tables **xiv**

Attestation of Authorship **xvi**

Acknowledgements **xix**

Copyright **xx**

Publications **xxi**

1 Introduction **1**

1.1 Problem Statement and Motivation 4
1.2 Research Questions 6
1.3 Research Design and Methodological Overview 7
1.4 Contributions and Significance 10
1.5 Included Publications and Their Role in the Thesis 12
1.6 Organisation of the Thesis 13
1.7 Chapter Conclusion 15

2 Prelude to the Systematic Literature Review **16**

2.1 Introduction 16
2.2 Role of the Review in the Thesis 17
2.3 Post-review Literature Update 17
2.4 Connection to the Following Chapters 20

3 Enhancing Emotional Well-Being With IoT Data Solutions for Depression: A Systematic Review **21**

3.1 Abstract 21

3.2 Introduction . . . 22
3.3 Systematic Literature Review Methodology . . . 24
3.3.1 Planning, Search, and Retrieval . . . 24
3.3.2 Screening and Selection . . . 24
3.3.3 Data Extraction and Analysis . . . 26
3.4 Data Collection Methods . . . 27
3.4.1 Sensor-based Data Collection . . . 27
3.4.2 User-generated and User-centric Data Collection . . . 30
3.4.3 Datasets . . . 31
3.5 Data Analysis Methods . . . 31
3.5.1 AI-based Methods . . . 31
3.5.2 Big Data . . . 33
3.5.3 Statistical Techniques . . . 33
3.5.4 Advanced Data Analysis Methods . . . 34
3.6 IoT-based Emotional Well-being . . . 35
3.6.1 Emotion Detection . . . 36
3.6.2 Emotion Analysis and Monitoring . . . 36
3.6.3 Therapeutic Approaches for Emotional Well-being . . . 38
3.7 IoT-based Solutions for Depression . . . 39
3.7.1 Predicting Depression . . . 40
3.7.2 Detecting Depression . . . 42
3.7.3 Managing Depression with IoT . . . 43
3.8 Discussion . . . 43
3.8.1 Data in Technological Emotional Well-being and Depression-Related Solutions . . . 43
3.8.2 Effectiveness Parameters in Emotional Well-being and Depression Solutions . . . 46
3.8.3 Depression Management . . . 47
3.9 Conclusion . . . 47

4 Prelude to Contextualising Environmental Data in IoT-Based Mental Health Monitoring 49
4.1 Introduction . . . 49
4.2 Role of Chapter 5 in the Thesis . . . 50
4.2.1 Interpretation and limitations of environmental stress labels . . . 51
4.3 Connection to Subsequent Chapters . . . 52

5 Integrating Environmental Data for Mental Health Monitoring: A Data-Driven IoT-Based Approach 53
5.1 Abstract . . . 53
5.2 Introduction . . . 54
5.2.1 Mental Health and Environmental Factors . . . 54
5.2.2 The Role of Big Data and IoT in Mental Health Management . . . 54
5.3 Related Work . . . 56
5.3.1 Environmental Factors and Mental Health . . . 56
5.3.2 Recent Advancements in Predictive Modelling . . . 56
5.3.3 Big Data in Healthcare . . . 57
5.3.4 Fuzzy Logic in Environmental Data Analysis . . . 57
5.4 Methodology . . . 57
5.4.1 Datasets . . . 58
5.4.2 Data Collection and Preprocessing . . . 58
5.4.3 Machine Learning and Big Data Integration . . . 60
5.4.4 System Architecture . . . 61
5.5 Results and Experiment . . . 63
5.5.1 Performance of Machine Learning Models . . . 63
5.5.2 Predictive Accuracy . . . 64
5.5.3 Class Imbalance and Model Performance . . . 64
5.6 Discussion . . . 67
5.6.1 Analysis of Model Performance . . . 68
5.6.2 Impact of Environmental Factors . . . 68
5.6.3 Limitations . . . 68
5.7 Conclusions . . . 69

6 Prelude to Gamified AI-Driven Depression Monitoring and Management 71
6.1 Introduction . . . 71
6.2 Role of Chapter 7 in the Thesis . . . 72
6.2.1 Synthetic Data Generation, Label Construction, and Evaluation Limitations . . . 73
6.3 Connection to Subsequent Chapters . . . 74

7 A Gamified AI-Driven System for Depression Monitoring and Management 75
7.1 Abstract . . . 75
7.2 Introduction . . . 76
7.3 Background and Related Works . . . 78

7.3.1 Gamification in Mental Health Tools . . . 78
7.3.2 Adaptive Daily Surveys for Emotional Monitoring . . . 78
7.3.3 AI for Personalized Depression Support . . . 79
7.3.4 Environmental Context: Temperature and Mood . . . 79
7.3.5 Motivation and Research Direction . . . 80
7.4 Materials and Methods: A Gamified AI-Driven System . . . 80
7.4.1 System Architecture . . . 81
7.4.2 System Component . . . 82
7.4.3 Embedded Research Capabilities . . . 84
7.4.4 Technical Implementation Details . . . 84
7.5 Results: Simulation-Based Evaluation of Depression Detection Models . . . 86
7.5.1 Evaluation Objectives . . . 86
7.5.2 Data Simulation Framework . . . 86
7.5.3 Machine Learning Pipeline . . . 88
7.5.4 Depression Scoring and Classification Algorithm . . . 93
7.5.5 Adaptive Survey Logic Simulation . . . 94
7.5.6 Ethical Considerations . . . 94
7.6 Discussion . . . 94
7.6.1 Validity of Synthetic Evaluation . . . 94
7.6.2 Adaptive and Context-Aware Modeling . . . 95
7.6.3 Research Utility and Generalizability . . . 95
7.7 Conclusions . . . 96

8 Prelude to the IoTMindCare Reference Architecture 98
8.1 Introduction . . . 98
8.2 Technical Context and Rationale . . . 98
8.3 Contribution of Chapter 9 . . . 99
8.4 Connection to Chapter 11 . . . 100

9 IoTMindCare: An Integrative Reference Architecture for Safe and Personalized IoT-Based Depression Management 101
9.1 Abstract . . . 101
9.2 Introduction . . . 102
9.3 Related Works . . . 104
9.3.1 Operational, Privacy, and Standards . . . 106
9.3.2 Synthesis: How the Literature Shaped IoTMindCare . . . 107
9.4 IoTMindCare: Proposed Reference Architecture . . . 107

9.4.1 Architectural Drivers . . . 107
9.4.2 Architecture Views . . . 110
9.5 Architecture Evaluation . . . 118
9.5.1 Comparison with Existing Architectures . . . 120
9.5.2 Scenario-Driven Instantiation and Configuration . . . 122
9.5.3 Limitations and Threats to Validity . . . 124
9.6 Conclusions . . . 125

10 Prelude to the Expert-Informed Refinement of IoTMindCare 127
10.1 Introduction . . . 127
10.2 Contribution to the Thesis . . . 128
10.3 Connection to Subsequent Chapters . . . 129

11 An Expert-Informed Architecture for Deployment-Oriented IoT-Based Depression Management 131
11.1 Abstract . . . 131
11.2 Introduction . . . 132
11.3 Literature Review . . . 134
11.3.1 IoT-based Behavioural Sensing for Mental Health . . . 134
11.3.2 Personalised Behavioural Modelling . . . 134
11.3.3 Machine Learning for Behavioural Anomaly Detection . . . 135
11.3.4 Alert Governance, Trust, and Workflow Integration . . . 135
11.3.5 Edge-Cloud Architectures in IoT Healthcare . . . 136
11.3.6 Regulatory and Interoperability Considerations . . . 136
11.3.7 Research Gaps . . . 137
11.4 Research Design . . . 137
11.5 Results . . . 142
11.5.1 Findings from Technical Expert Feedback . . . 142
11.5.2 Findings from Clinical Expert Feedback . . . 147
11.5.3 Extending the IoTMindCare Architecture . . . 148
11.6 Discussion . . . 150
11.6.1 Threats to Validity . . . 153
11.7 Conclusions and Future Work . . . 154

12 Discussion 156
12.1 Introduction . . . 156
12.2 Addressing the Research Questions . . . 158

12.2.1 Overarching Research Question . . . 158
12.2.2 RQ1: Limitations and Design Requirements in Existing IoT-Based Depression Systems . . . 159
12.2.3 RQ2: Contextual and Behavioural Data for Personalised and Meaningful Monitoring . . . 160
12.2.4 RQ3: Organisation of IoT Components within an Integrative Reference Architecture . . . 161
12.2.5 RQ4: Expert-Informed Refinement of the Architecture . . . 162
12.3 Synthesis of Thesis-Level Findings . . . 164
12.3.1 From Monitoring to Management . . . 165
12.3.2 From Generic Prediction to Personalisation . . . 166
12.3.3 From Single-Source Data to Multimodal and Context-Aware Interpretation . . . 167
12.3.4 From Prototype Functionality to Architecture . . . 168
12.3.5 From Automated Decision-Making to Governed Human-in-the-Loop Support . . . 169
12.3.6 From Technical Possibility to Implementation-Oriented Design . . . 171
12.4 Overall Contribution of the Thesis . . . 172
12.5 Limitations and Threats to Validity . . . 173
12.5.1 Construct Validity . . . 174
12.5.2 Internal Validity . . . 175
12.5.3 External Validity and Transferability . . . 175
12.5.4 Conclusion Validity . . . 176
12.5.5 Reliability and Dependability . . . 176
12.6 Future Work . . . 177
12.6.1 Prototype Implementation of IoTMindCare . . . 177
12.6.2 Simulation-Based Evaluation of Safety and Alert Governance . . . 178
12.6.3 User-Centred Evaluation with Key Stakeholders . . . 178
12.6.4 Longitudinal Pilot Evaluation . . . 179
12.6.5 Clinical and Organisational Validation . . . 180
12.6.6 Cross-Cutting Future Research Directions . . . 180
12.7 Chapter Conclusion . . . 181

13 Conclusion 183
13.1 Introduction . . . 183
13.2 Summary of the Thesis . . . 183
13.3 Overall Contribution . . . 184

13.4 Final Concluding Remarks 185

References **186**

Appendix A Ethics Approval **206**

Appendix B Participant Information Sheet **208**

Appendix C Questionnaire **212**

Appendix D Participant Feedback Reports **229**

List of figures

1.1 Trends in the estimated prevalence of depressive disorders globally and in New Zealand, expressed as a percentage of the population, based on Global Burden of Disease 2023 estimates [1]. 2
1.2 Overall research pathway and contribution of the thesis. 7

3.1 Publications per year for the provided search string 25
3.2 Systematic Literature Review (SLR) process 26
3.3 Number of data collection methods (blue: sensor-based methods, orange: user-centric methods, green: datasets) 27
3.4 Data analysis methods in emotional well-being and depression healthcare systems 32
3.5 Trend of using AI for analysing data in emotional well-being and depression healthcare solutions including surveys 32
3.6 Emotional well-being solutions based on types and methods 35
3.7 Depression solutions based on type and method 40
3.8 Data in IoT-based emotional well-being and depression systems 44

5.1 System Architecture: A comprehensive framework integrating IoT devices, data processing platforms, and machine learning models for real-time stress monitoring and prediction. 62
5.2 Comparison of Precision, Recall, and F1 Score Across Machine Learning Models: SVM Outperforms LR and LSTM in Handling Static Environmental Data. 64
5.3 Confusion Matrix for Logistic Regression (LR): The model performs well for dominant classes but struggles with mid-range categories. 65
5.4 Confusion Matrix for Support Vector Machine (SVM): The model shows improved accuracy for “Very High” and slight improvements for “Low” and “Medium” classes. 66

5.5 Confusion Matrix for Long Short-Term Memory (LSTM): The model excels in predicting "Very High" but fails entirely for "Low", "Medium", and "High" classes. 66

7.1 Physical view of the system architecture showing mobile front-end, AWS back-end integration, and AI analytics components. 83
7.2 Sample of stored user data including mood rating, location, and gameplay interaction timestamps. 84
7.3 Main application interface showing game selection and real-time weather contextualization, and the game page (Sudoku). 85
7.4 Feature importance based on permutation analysis. Mood score, game accuracy, and temperature are among the top contributors to prediction performance. This highlights the value of integrating behavioral, cognitive, and environmental data. 90
7.5 ROC curves for all three models. The neural network shows the highest AUC, indicating strong discriminative performance. Random forest also performs well across all thresholds. Logistic regression has lower sensitivity, especially at lower thresholds. 91
7.6 Confusion matrices for each classifier, showing predicted vs. actual labels. Random forest and neural network both achieve high accuracy, with minimal false negatives. Logistic regression shows more misclassifications, particularly in the depressed class. 92

9.1 Graphical abstract of the main contributions of *IoTMindCare* 105
9.2 Development view - UML Component diagram showing the major software modules for smart home, cloud, and user environments 114
9.3 Process view - UML Sequence Diagram showing end-to-end system behavior categorized into four stages and featuring dynamic interactions among smart home, cloud, and user interfaces 116
9.4 Deployment view - UML Deployment diagram showing the mapping of the development view components to physical or infrastructure resources . . . 117

11.1 Workflow of the mixed-method expert evaluation, showing the initial architecture context, expert recruitment, questionnaire design, data collection, analysis, and architecture refinement. 140
11.2 Logical architecture view of IoTMindCare, showing the functional layering of the system and the cross-cutting roles of progressive personalisation, safety and clinical oversight, and governance and interoperability. 149

11.3 Hybrid edge-cloud processing view of IoTMindCare, showing the allocation of real-time screening, local resilience, and immediate alerting to edge resources, and the allocation of longitudinal analytics, model adaptation, interoperability services, and governance support to cloud resources. 150

12.1 The six key transitions offered through *IoTMindCare*. 164

List of tables

1.1 Research contributions, evidence, and implementation boundaries across the thesis. . . . 14

2.1 Summary of post-review literature developments identified from 13 directly relevant studies published after 15 May 2024 . . . 18

3.1 Comparison of the coverage of our research with existing surveys . . . 23
3.2 Inclusion and exclusion criteria . . . 25
3.3 Data collection via wearables, categorised by type of data, wearable type, and purpose of data collection . . . 28
3.4 Data collection by IoT sensors, categorised by type of data, sensor type, and purpose of data collection . . . 29
3.5 Data collection by questionnaires and surveys, categorised by question types or nature, and their aim of collecting data . . . 30
3.6 Data collection by mobile applications, categorised by type of data and method of collection . . . 31
3.7 AI-based solutions categorised by type of algorithm and purpose of analysis 33
3.8 Methods for analysing and monitoring emotions . . . 37
3.9 Capabilities of depression-related IoT solutions . . . 41
3.10 Types of data used in technological emotional well-being and depression systems . . . 45

5.1 Categorisation of Stress Levels Based on Temperature and Humidity. . . . 59
5.2 Performance Metrics of Machine Learning Models for Stress Level Classification. . . . 63

7.1 Structure of the synthetic dataset categorized by feature, data type, and generation method. . . . 88
7.2 Model performance on simulated depression dataset. . . . 89

9.1 Comparison of recent IoT-based mental and physical health monitoring architectures, their characteristics, and associated benefits. 108
9.2 Scenario for Requirement **Safety** 111
9.3 Scenario for Requirement **Personalization** 112
9.4 Logical view - layering diagram showing the five architectural layers of the architecture and their mapping to safety and personalization 113
9.5 Mapping Quality Attribute Scenarios to architectural evidence 119
9.6 Comparison with recent IoT-health architectures 121
9.7 Scenario to architecture mapping: signals, literature sources, relevant views, primary QAS and practical configuration notes for prototype instantiation. . 123
9.8 List of Abbreviations 126

11.1 Representative studies on IoT-based behavioural monitoring for mental health in relation to key architectural dimensions 138
11.2 Technical Expert Feedback: Calibration, personalised initialisation, and generic model setup 143
11.3 Deployment architecture, cloud platform selection, and cloud integration . . 144
11.4 Anomaly detection methods and regulatory/standards considerations 146
11.5 Main themes from open-ended responses by clinical experts 148
11.6 Traceable mapping from expert findings to architectural refinements 151

12.1 Summary of thesis-level threats to validity. 174

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 used artificial intelligence tools or generative artificial intelligence tools (unless it is clearly stated, and referenced, along with the purpose of use), 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.

Sanaz Zamani
October 2026

Co-Authored Works and Declaration of Collaboration

STUDENT AND SUPERVISOR APPROVALS

By signing you are confirming that the co-author contributions stated in the table(s) below are accurate.

Student Name	Sanaz Zamani	**Signature**		**Date**	08/07/2026
Supervisor Name	Minh Nguyen	**Signature**		**Date**	09/07/2026
Supervisor Name	Samaneh (Sam) Madanian	**Signature**		**Date**	09/07/2026
Supervisor Name	Roopak Sinha	**Signature**		**Date**	08/07/2026

Chapter Number:	**Chapter 3**
Manuscript Title:	**Enhancing Emotional Well-being with IoT Data Solutions for Depression: A Systematic Review**
Publication Status:	**Published**
Reference if published:	**Zamani, S., Sinha, R., Nguyen, M., & Madanian, S. (2025). Enhancing emotional well-being with IoT data solutions for depression: A systematic review. *IEEE Journal of Biomedical and Health Informatics*, *29*(3), 1919-1930.**

AUTHOR SURNAME: (order as per manuscript)	CONTRIBUTION (May copy from the guidelines above)
Zamani	Study design, literature search, screening, analysis, interpretation, and manuscript drafting.
Sinha	Contributed technical and architectural guidance and critically revised the manuscript.
Nguyen	Provided supervisory guidance and critically revised the manuscript.
Madanian	Contributed health informatics expertise and critically revised the manuscript.

Chapter Number:	**Chapter 5**
Manuscript Title:	**Integrating Environmental Data for Mental Health Monitoring: A Data-Driven IoT-Based Approach**
Publication Status:	**Published**
Reference if published:	**Zamani, S., Nguyen, M., & Sinha, R. (2025). Integrating environmental data for mental health monitoring: A data-driven IoT-based approach. *Applied Sciences*, *15*(2), 912.**

AUTHOR SURNAME: (order as per manuscript)	CONTRIBUTION (May copy from the guidelines above)
Zamani	Study design, data preparation, analysis, interpretation, and manuscript drafting.
Nguyen	Provided methodological and supervisory guidance and critically revised the manuscript.
Sinha	Contributed IoT and technical guidance and critically revised the manuscript.

Chapter Number:	**Chapter 7**
Manuscript Title:	**A Gamified AI-Driven System for Depression Monitoring and Management**
Publication Status:	**Published**
Reference if published:	**Zamani, S., Rostami, A., Nguyen, M., Sinha, R., & Madanian, S. (2025). A gamified AI-driven system for depression monitoring and management. *Applied Sciences*, *15*(13), 7088.**

AUTHOR SURNAME: (order as per manuscript)	CONTRIBUTION (May copy from the guidelines above)
Zamani	System design, analysis, interpretation, and manuscript drafting.
Rostami	Contributed to the system design and interpretation.
Nguyen	Provided supervisory guidance and critically revised the manuscript.
Sinha	Contributed IoT and system-design expertise and critically revised the manuscript.
Madanian	Contributed health informatics expertise and critically revised the manuscript.

Chapter Number:	**Chapter 9**
Manuscript Title:	**IoTMindCare: An Integrative Reference Architecture for Safe and Personalized IoT-Based Depression**
Publication Status:	**Published**
Reference if published:	**Zamani, S., Sinha, R., Madanian, S., & Nguyen, M. (2025). IoTMindCare: An Integrative Reference Architecture for Safe and Personalized IoT-Based Depression Management. *Sensors*, *25*(22), 6994.**

AUTHOR SURNAME: (order as per manuscript)	CONTRIBUTION (May copy from the guidelines above)
Zamani	Architecture design, analysis, interpretation, and manuscript drafting.
Sinha	Contributed architectural, IoT, and system design and critically revised the manuscript.
Nguyen	Provided methodological and supervisory guidance and critically revised the manuscript.
Madanian	Contributed health informatics expertise and critically revised the manuscript.

Chapter Number:	**Chapter 11**
Manuscript Title:	**An Expert-Informed Architecture for Deployment-Oriented IoT-Based Depression Management**
Publication Status:	**Submitted for Publication**
Reference if published:	**Zamani, S., Sinha, R., Madanian, S., & Nguyen, M. (2026). An expert-informed architecture for deployment-oriented IoT-based depression management. Submitted to Journal of Healthcare Informatics Research.**

AUTHOR SURNAME: (order as per manuscript)	CONTRIBUTION (May copy from the guidelines above)
Zamani	Study design, data collection, analysis, interpretation, architecture refinement, and manuscript drafting.
Sinha	Contributed technical and architectural expertise and critically revised the manuscript.
Madanian	Contributed health informatics and ethics expertise, and critically revised the manuscript.
Nguyen	Provided methodological and supervisory guidance and critically revised the manuscript.

Acknowledgements

I would like to express my sincere gratitude to everyone who supported me throughout my PhD journey.

First and foremost, I would like to express my deepest appreciation to my supervisors, Prof. Minh Nguyen, Prof. Samaneh Madanian, and Prof. Roopak Sinha, for their guidance, encouragement, and support throughout this research. I am especially grateful to Prof. Roopak Sinha for his continuous support, detailed feedback, patience, and generous guidance during the most challenging stages of this thesis. His thoughtful advice and belief in my work played an important role in helping me complete this research.

I also sincerely thank Prof. Minh Nguyen and Prof. Samaneh Madanian for their valuable academic guidance, constructive feedback, and support. Their expertise and insights contributed significantly to the development, refinement, and completion of this thesis.

I am grateful to Auckland University of Technology for providing the academic environment, resources, and support that enabled me to undertake this research.

Ethics approval for the human-participant component of this research was granted by the Auckland University of Technology Ethics Committee (AUTEC), application number 25/401, on 9 December 2025.

My sincere appreciation goes to the experts who participated in this research. Their time, professional insights, and valuable feedback made an important contribution to the expert-informed refinement of the proposed architecture.

Finally, I would like to express my deepest gratitude to my family for their unconditional love, patience, and support. I am especially thankful to my parents for their endless love, sacrifices, encouragement, and belief in me throughout my life and academic journey. I am also deeply grateful to my husband for his love, patience, understanding, and unwavering support during the most demanding stages of my PhD. This achievement would not have been possible without the love, sacrifices, and support of my family.

Copyright

Theses, dissertations and research projects are protected by the Copyright Act 1994 (New Zealand). This thesis, dissertation or research projects may be consulted by you, provided you comply with the provisions of the Act and the following conditions of use:

- Any use you make of these documents or images must be for research or private study purposes only, and you may not make them available to any other person.
- Authors control the copyright of their thesis, dissertation or research project. You will recognise the author's right to be identified as the author of the thesis, dissertation or research project, and due acknowledgment will be made to the author where appropriate.
- You will obtain the author's permission before publishing any material from the thesis, dissertation or research project.
- The ownership of any intellectual property rights which may be described in this thesis is vested in the Auckland University of Technology, subject to any prior agreement to the contrary, and may not be made available for use by third parties without the written permission of the University, which will prescribe the terms and conditions of any such agreement.

Copyright @2026. Sanaz Zamani

Publications

Zamani, S., Sinha, R., Nguyen, M., & Madanian, S. (2025). Enhancing emotional well-being with IoT data solutions for depression: A systematic review. IEEE Journal of Biomedical and Health Informatics, 29(3), 1919-1930.

Zamani, S., Nguyen, M., & Sinha, R. (2025). Integrating environmental data for mental health monitoring: A data-driven IoT-based approach. Applied Sciences, 15(2), 912.

Zamani, S., Rostami, A., Nguyen, M., Sinha, R., & Madanian, S. (2025). A gamified AI-driven system for depression monitoring and management. Applied Sciences, 15(13), 7088.

Zamani, S., Sinha, R., Madanian, S., & Nguyen, M. (2025). IoTMindCare: An Integrative Reference Architecture for Safe and Personalized IoT-Based Depression Management. Sensors, 25(22), 6994.

Zamani, S., Sinha, R., Nguyen, M., & Madanian, S. (2026). A Multimodal AI–IoT Framework for Monitoring Student Well-Being in Hybrid and Transnational Learning Environments. In 2026 Australasian Computer Science Week (pp. 60-63).

Zamani, S., Sinha, R., Madanian, S., & Nguyen, M. (2026). An expert-informed architecture for deployment-oriented IoT-based depression management. Submitted to Journal of Healthcare Informatics Research.

Chapter 1

Introduction

Depression is a complex mental health condition that affects emotional well-being, daily functioning, relationships, work, and study. The World Health Organization (WHO) describes depression as a common mental disorder, and based on 2021 Global Burden of Disease estimates, reports that approximately 332 million people worldwide are living with depression [1, 2]. It also notes that treatment gaps remain substantial, including in high-income countries [2]. More broadly, the WHO's World Mental Health Report argues that many people who need mental health support do not receive adequate care, and that current mental health systems remain under-resourced, fragmented, and insufficiently responsive to population needs [3]. Consequently, the prevailing view is to consider depression not only as a clinical issue, but also as a systems-level public health challenge that requires better models of detection, monitoring, support, and continuity of care [4, 5].

In Aotearoa New Zealand, recent population-level data provide local context for the importance of improving approaches to depression-related monitoring and support. The 2024/25 New Zealand Health Survey reported that 14.3% of adults aged 15 years and over experienced high or very high psychological distress in the four weeks before the survey, compared with 7.4% in 2019/20 [6]. Although psychological distress is broader than depression and should not be interpreted as a depression-specific diagnosis, this increase indicates a growing level of emotional distress in the population and reinforces the need to explore safe, personalised, and context-aware approaches that may support earlier awareness and more continuous monitoring.

To further illustrate the scale and persistence of depression as a public health challenge, Figure 1.1 presents percentage-based trends in the estimated prevalence of depressive disorders globally and in New Zealand from 1990 to 2023, based on Global Burden of Disease 2023 estimates [1]. The figure provides contextual motivation for the thesis by showing that depression remains a sustained population-level concern at both global and local levels.

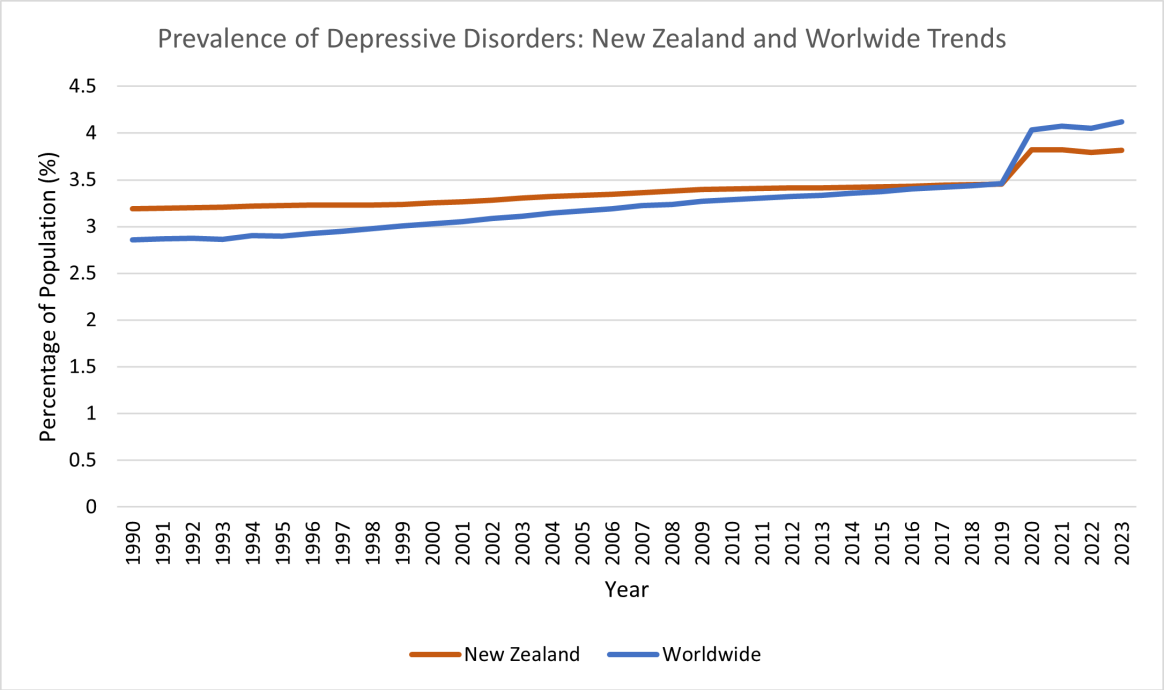


Fig. 1.1 Trends in the estimated prevalence of depressive disorders globally and in New Zealand, expressed as a percentage of the population, based on Global Burden of Disease 2023 estimates [1].

Traditional approaches to depression management remain heavily dependent on periodic clinical encounters and self-reported symptoms [7]. Standardised instruments such as symptom questionnaires play an important role in assessment, screening, and monitoring, but they are usually administered at discrete points in time [8]. As a result, they provide limited visibility into everyday fluctuations, behavioural changes, contextual triggers, and gradual deterioration that may occur between formal assessments. This limitation is particularly important for depression because symptoms often unfold over time and may be reflected in changes to sleep, physical activity, social interaction, daily routine, motivation, and participation in everyday activities. Therefore, a major challenge for depression management is how to support a more continuous, personalised, and context-aware understanding of an individual's condition without replacing clinical judgement or increasing the burden on users and care providers.

Digital health has opened new possibilities for delivering health services beyond episodic, clinic-based encounters. The WHO guidelines on digital interventions for health system strengthening stress that digital health tools should be assessed for benefits, harms, acceptability, feasibility, resource use, and equity, rather than being treated as inherently beneficial simply because they are technological [9]. In mental health specifically, recent studies empha-

sise both the promise and the complexity of digital innovation, noting that implementation, ethics, regulation, workflow fit, and service integration remain central challenges [10, 11]. This broader digital health context is important because it frames the development of Internet of Things (IoT)-based mental health systems as a health-systems and architecture design problem, not just a software, sensor, or machine learning problem.

In the context of IoT-enabled mental health monitoring, smartphones, wearables, and ambient sensors have created opportunities for continuous and multimodal observation of behaviour and context in daily life. In mental health research, data from smartphones and other personal digital devices can be used to continuously measure behaviour, activity, and other health-related patterns in real-world settings [12]. This scope is especially relevant to depression because depressive symptoms are often reflected in changes in sleep, activity, routine, mobility, social behaviour, and device use over time rather than through a single static signal [13]. The potential of IoT-enabled systems lies in their ability to combine heterogeneous streams of behavioural, physiological, and environmental information to support a more timely and ecologically valid understanding of mental health states [14].

The availability of sensing technologies does not by itself solve the problem of depression management. A systematic review of smartphone-based digital sensing methods in depression found growing interest in these approaches, but also concluded that studies had achieved only moderate performance and that the field still faced substantial methodological and translational limitations [13]. Likewise, a scoping review of passive sensing with wearables and smartphones reported that relatively small cohort studies dominated the evidence base and that most work focused on depression and anxiety using a limited set of biomarkers [15]. These findings suggest that the field is moving, but not yet mature enough to assume that data collection alone will translate into robust clinical or operational value.

A further challenge is that mental health data are inherently contextual [16]. Passive sensing studies have highlighted physiological and behavioural variables, including movement, sleep, heart rate, voice features, social activity proxies, and device use. However, the meaning of these variables is rarely universal across people or situations. A reduction in activity, for example, may indicate depressive withdrawal for one person, recovery from illness for another, a planned rest day for another, or an artefact caused by missing data or device non-use [17]. Sheikh et al. [14] review wearable, smartphone-based, and environmental sensing and explicitly note the need to consider both intra-personal and inter-personal differences when interpreting passive data. This means that context-awareness and personalisation are not optional enhancements to IoT-based depression monitoring; they are central requirements for valid interpretation.

In addition to contextual interpretation and personalisation, sustained user engagement is also essential. Digital mental health interventions may expand access, convenience, and flexibility, and they are often acceptable to users [18]. For example, Lau et al. [18] found generally positive acceptability among patients with depression or anxiety disorders, although the quality of the evidence was only moderate. However, acceptability is not the same as sustained engagement. Smith et al. [19] concluded that engagement rates in digital mental health are consistently low and identified three major challenges: the lack of agreed-upon engagement metrics, the lack of evidence linking engagement quality to outcomes, and the lack of standards for user involvement. These observations are crucial for a thesis on depression management because a technically capable system that users do not meaningfully engage with is unlikely to produce lasting benefit.

Extant literature supports a clear conclusion: the field no longer needs only more sensing studies or isolated prediction models. Further progress also requires architecture-level thinking about how sensing, interpretation, safety, personalisation, engagement, interoperability, stakeholder interaction, and governance can be organised within coherent systems. That is the context in which this thesis is positioned. The thesis investigates how fragmented work in IoT-based depression monitoring and management can be consolidated into an integrative reference architecture that supports safe, personalised, and context-aware system design, while providing a foundation for future implementation and validation.

1.1 Problem Statement and Motivation

The systematic review presented in Chapter 3 shows that, although IoT-enabled and digitally mediated mental health research has grown rapidly, current work on depression remains fragmented across several design dimensions [20]. Four dimensions are particularly relevant to the architecture-level contribution of this thesis:

1. Much of the literature is component-focused rather than system-focused. Many studies examine whether a particular sensor, feature set, or machine learning model can predict depressive symptoms, but fewer studies address how sensing, data management, interpretation, user feedback, alerting, escalation, and governance should be organised into a coherent end-to-end system [10]. Recent field-level reviews have also noted that digital mental health innovation is progressing alongside unresolved questions about implementation, evidence quality, ethics, workflow relevance, and service transformation [11]. This indicates the need for architecture-level thinking that clarifies how different system responsibilities should fit together.

2. Safety and personalisation require stronger architectural attention. In depression management, outputs that are poorly calibrated, context-insensitive, or overly generic may lead to false reassurance, false alerts, alert fatigue, unnecessary anxiety, or user distrust. At the same time, the current sensing literature makes clear that mental health signals vary across individuals and over time, making one-size-fits-all interpretation problematic [13]. Therefore, the challenge is not simply to gather more data, but to design systems that can accommodate individual baselines while remaining clinically and ethically cautious [14, 10].

3. Contextual interpretation remains insufficiently integrated into depression-oriented system design. Environmental and situational factors can shape behaviour, routine, comfort, sleep, and emotional experience, yet they are still underused in many monitoring approaches. The passive sensing literature shows that environmental sensing is part of the broader digital sensing ecosystem, but it has not been consistently incorporated into depression-oriented system design [14]. This creates a gap between what can technically be sensed and what is meaningfully interpreted in context. An architecture for depression monitoring should provide a way to relate behavioural variation to contextual, environmental, and data-quality factors before treating it as clinically meaningful.

4. Sustained engagement is an important design consideration for depression monitoring and management. Evidence suggests that digital mental health interventions can be acceptable, but engagement and retention are often inconsistent [18]. This is particularly important in depression, where reduced motivation and energy may themselves affect continued use [19]. This thesis does not claim to solve the problem of long-term engagement empirically. Rather, it treats engagement as an architectural and design requirement by considering how user interaction, feedback, reflection, and self-management can be supported within IoT-based depression management systems [21].

The central problem addressed by this research can be stated as follows:

Current IoT-based approaches to depression monitoring and management remain fragmented and insufficiently integrated at the architecture level. Existing work often emphasises isolated sensing, prediction, or prototype functionality, while giving less systematic attention to how sensing, interpretation, safety, personalisation, contextual awareness, engagement, stakeholder interaction, interoperability, and governance should be organised within a coherent system architecture.

This thesis is motivated by the need to move beyond isolated technical demonstrations of connected sensing devices, smartphone-based data collection, mobile applications, and machine learning models toward a coherent architectural foundation for future IoT-enabled depression monitoring and management systems. In this thesis, IoT-enabled systems are understood as systems that combine connected sensing devices, data communication, data processing, analytics, feedback, and stakeholder-facing interfaces within an integrated architecture. The thesis develops and refines an architectural foundation that can inform future implementation, simulation, pilot studies, and empirical validation.

Accordingly, the problem addressed in this thesis is the absence of a coherent architectural foundation that explains how multiple IoT components should be integrated, governed, personalised, and evaluated for safe and context-aware depression monitoring and management.

1.2 Research Questions

The research questions in this thesis are framed around a single doctoral problem: how to consolidate fragmented IoT-based depression monitoring and management research into a coherent reference architecture that supports safe, personalised, context-aware, and future-adaptable system design. The questions guide the research pathway through which *IoTMindCare* emerges as the integrated architectural outcome of the thesis.

The overarching research question is:

> *How can architecture-level requirements and design principles be derived and organised to address the fragmentation of current IoT-enabled depression monitoring and management approaches across sensing, interpretation, safety, personalisation, contextual awareness, engagement, stakeholder interaction, interoperability, and governance?*

This overarching question is addressed through four integrated research questions:

- **RQ1:** What limitations and design requirements characterise existing IoT-based approaches to depression monitoring and emotional well-being?
- **RQ2:** How can the limitations identified in existing IoT-based approaches concerning environmental contextualisation and user engagement be explored to inform the design of personalised and meaningful depression monitoring and management systems?
- **RQ3:** How can the literature-derived requirements and design insights from contextual sensing and engagement-oriented studies be integrated into a reference architecture

that addresses safety, personalisation, stakeholder interaction, interoperability, and governance for IoT-based depression monitoring and management?

- **RQ4:** How can technical and clinical expert perspectives inform the refinement of the proposed reference architecture to strengthen its technical and clinical relevance, safety considerations, and suitability for future implementation and validation?

Figure 1.2 summarises the overall research pathway and contribution of the thesis. It shows how the prelude chapters, publication-based chapters, and research questions connect across the thesis and how the individual studies collectively lead to the development and refinement of *IoTMindCare* as the overall architectural contribution.

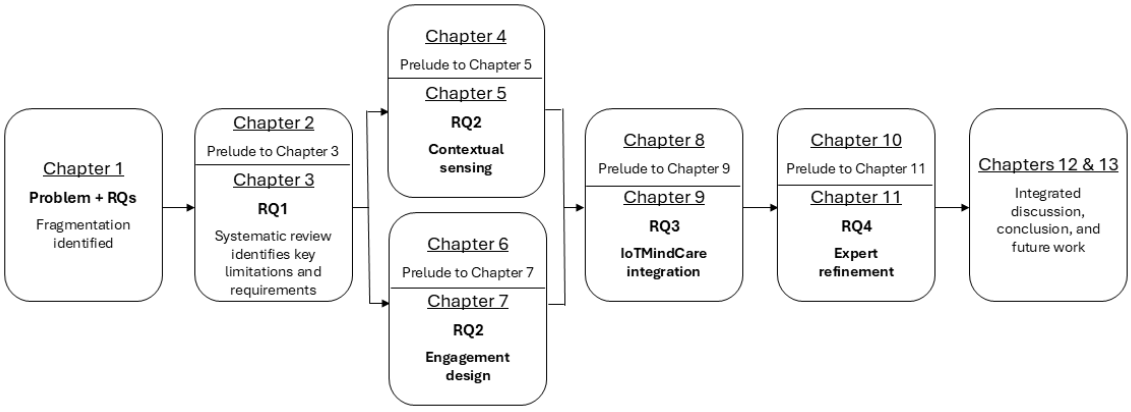


Fig. 1.2 Overall research pathway and contribution of the thesis.

As shown in Figure 1.2, the thesis progresses from problem formulation and evidence synthesis to design exploration, architectural integration, and expert-informed refinement. This structure reinforces the position that *IoTMindCare* is not assumed at the outset, but emerges as the integrated architectural outcome of the research questions and publication-based studies.

1.3 Research Design and Methodological Overview

This thesis adopts a thesis-by-publication format and is guided by a Design Science Research methodology [22]. Design Science Research is appropriate for information systems and engineering research in which the purpose is to develop and evaluate artefacts that address identified problems [22, 23]. In this thesis, the artefact is an architecture-level contribution: the development and refinement of an integrative reference architecture for safe, personalised, and context-aware IoT-enabled depression monitoring and management.

In this thesis, architecture-level evaluation refers to the assessment of the proposed reference architecture through architectural modelling, quality attribute scenarios, ATAM-informed trade-off analysis, comparison with existing architectures, and expert-informed refinement. These methods provide evidence of architectural coherence, requirements coverage, design rationale, and perceived technical and clinical relevance. Algorithmic evaluation is limited to the specific data-driven studies presented in Chapters 5 and 7. IoTMindCare is presented as an evidence-informed and expert-refined reference architecture intended to guide future implementation and validation (Chapters 9 and 11).

Design Science Research is suitable for this thesis because it addresses how fragmented technical, human, contextual, safety, and clinical requirements can be organised into a coherent architecture for future IoT-enabled depression management systems. Therefore, the research follows a design-oriented pathway in which the problem is investigated, design requirements are identified, specific design dimensions are explored, an architectural artefact is developed, and the resulting architecture is refined through expert-informed refinement.

The research begins with a systematic review of IoT-based approaches for depression and emotional well-being. This stage addresses **RQ1** by establishing the state of the field and identifying limitations and design requirements in existing IoT-based approaches. In particular, the review examines limitations across sensing approaches, analytics, user engagement, privacy, security, interoperability, system integration, and practical implementation considerations. It provides the evidence base for the thesis by showing that existing work is fragmented and that an architecture-level response is needed.

The thesis then examines contextual sensing through a data-driven IoT-based approach that integrates environmental data into mental health monitoring. This stage contributes to **RQ2** by exploring how contextual data can inform more personalised and meaningful interpretation of mental health-related patterns. The study focuses on environmental conditions, particularly temperature and humidity, as contextual variables that may influence or help interpret behaviour and well-being. The purpose is not to claim that environmental data alone can diagnose depression, but to examine how environmental context can be incorporated into IoT-enabled monitoring and personalised interpretation.

The next stage examines engagement and self-management through a gamified AI-driven system for depression monitoring and management. This stage also contributes to **RQ2** by exploring how user interaction and engagement-oriented design can inform IoT-enabled depression monitoring and management. The study considers how micro-games, adaptive surveys, sensor data, personalised feedback, and AI analytics may support user interaction, reflection, and self-awareness. It does not claim to solve the problem of sustained engagement

empirically; rather, it treats engagement as a design requirement that should be considered within future IoT-enabled depression management systems.

Building on the systematic review, contextual sensing study, and engagement-oriented system design, the thesis then develops *IoTMindCare* as an integrative reference architecture. This stage addresses **RQ3** by organising safety, personalisation, stakeholder interaction, and governance within a coherent architecture for IoT-enabled depression monitoring and management. The architecture brings together multimodal sensing, hybrid edge-cloud processing, personalised analytics, anomaly detection, feedback mechanisms, stakeholder-centred design, and safety considerations. This stage shifts the thesis from individual design dimensions to an integrated architectural outcome.

Finally, the thesis refines *IoTMindCare* through an expert-informed evaluation that integrates technical and clinical perspectives. This stage addresses **RQ4** by examining how expert feedback can inform the refinement of the resulting architecture and indicate its architectural plausibility, perceived usefulness, technical relevance, and perceived clinical relevance. The evaluation considers issues such as personalisation, alert governance, clinical workflow relevance, interoperability, privacy, and practical implementation considerations. The expert-informed study refines the architecture by addressing technical feasibility considerations, perceived clinical relevance, and practical implementation requirements, providing a more informed foundation for future empirical evaluation. As this stage involved expert participants, the related ethics approval, participant information sheet, questionnaire, and participant feedback reports are provided in Appendices A, B, C, and D, respectively.

Overall, the methodological pathway can be summarised as follows:

> *Problem investigation and requirements identification (RQ1) → design exploration through contextual sensing and engagement-oriented design (RQ2) → architecture artefact construction (RQ3) → expert-informed evaluation and refinement (RQ4).*

This pathway allows the thesis to build a coherent contribution while recognising the complexity of depression management and the need for both technical and clinically informed design. This combination of methods is appropriate because the research problem is both technical and socio-clinical. A systematic review was needed to establish the evidence base and identify gaps. Data-driven and gamified system studies were needed to explore specific design requirements related to context and engagement. Architecture development was needed to integrate these requirements into a coherent system structure. An expert-informed evaluation was needed to obtain professional judgements on the architecture's perceived technical feasibility, clinical relevance, and implementation requirements. Together, these

methods support the development and refinement of a reference architecture that is evidence-informed, design-oriented, clinically cautious, and intended to inform future implementation, evaluation, and validation.

1.4 Contributions and Significance

The main contribution of this thesis is the development and refinement of *IoTMindCare* as an integrative reference architecture for safe, personalised, and context-aware IoT-based depression monitoring and management. The thesis responds to a fragmented research field in which many existing approaches focus on isolated sensing, prediction, or prototype-level functionality, while giving less attention to how these components should be integrated, interpreted, personalised, governed, and evaluated within a coherent architecture.

The thesis makes five interrelated contributions to the design of safe, personalised, context-aware, and architecture-centred IoT-based depression management:

1. The thesis provides a system-level understanding of the limitations of current IoT-based systems for depression and emotional well-being [20]. Rather than treating the field as a collection of separate technologies, the thesis identifies fragmentation across sensing, analytics, engagement, privacy, interoperability, system integration, safety, and personalisation. This contribution establishes the need for an architecture-level response that can organise these concerns into a coherent framework.

2. The thesis contributes a context-aware interpretation perspective for IoT-based depression monitoring. It shows that mental health-related data should not be interpreted in isolation, but in relation to the user's environment, behaviour, routine, and situation. By incorporating environmental variables such as temperature and humidity into the monitoring process, the thesis strengthens the argument that contextual information is necessary for more meaningful and personalised interpretation of depression-related patterns [24].

3. The thesis contributes an engagement-oriented perspective on IoT-based depression management. It argues that passive monitoring alone may be insufficient for meaningful depression support because user participation, reflection, feedback, and self-management are important design considerations. Through the exploration of gamified AI-IoT mechanisms, the thesis examines how user interaction, adaptive surveys, micro-games, and personalised feedback can inform the design of more engaging and participatory depression management systems [25].

4. The thesis contributes *IoTMindCare* as an integrative reference architecture that brings together the technical, human, contextual, safety, and personalisation requirements of IoT-based depression monitoring and management [26]. The architecture organises multimodal sensing, communication, edge-cloud processing, anomaly detection and safety considerations, personalised analytics, feedback mechanisms, stakeholder roles, and interoperability into a coherent framework. This contribution moves the research beyond individual system components and provides architectural guidance for future IoT-enabled depression management systems.

5. The thesis contributes an expert-informed refinement of *IoTMindCare*, strengthening the consideration of technical feasibility, safety requirements, perceived clinical relevance, and practical implementation constraints within the proposed design. By integrating technical and clinical perspectives, the thesis identifies and addresses key concerns across calibration, alert governance, workflow relevance, privacy, interoperability, clinical oversight, and future implementation considerations. This contribution is important because IoT-based depression management systems must not only be technically feasible but also safe, interpretable, acceptable, and aligned with clinical and practical settings.

The significance of the thesis lies in its shift from technology-centred monitoring toward architecture-centred depression management. Rather than presenting IoT-based depression monitoring as a purely predictive task, the thesis demonstrates that safe and meaningful depression management requires contextual interpretation, personalised calibration, user engagement, human-in-the-loop oversight, alert governance, privacy-aware design, interoperability, and implementation-informed system organisation. Therefore, *IoTMindCare* provides a foundation that future researchers and developers may use to guide prototyping, simulation, digital twin modelling, pilot studies, longitudinal evaluation, and clinical validation of IoT-enabled systems for depression monitoring and management.

This thesis's significance is in providing a coherent architectural foundation through which future systems can be designed, implemented, evaluated, and refined in a safer and more personalised manner. In doing so, the thesis contributes to the broader field of digital mental health by showing how IoT-based depression management can move beyond fragmented prototypes toward more integrated, responsible, and architecture-guided system design.

1.5 Included Publications and Their Role in the Thesis

This thesis includes five publications or manuscripts that together form a coherent research programme. Although each publication addresses a specific research problem, their combined contribution is the development and refinement of *IoTMindCare*. The publications should be read not as separate studies, but as connected stages in a progressive design pathway.

1. **Zamani, S., Sinha, R., Nguyen, M., & Madanian, S. (2025). Enhancing emotional well-being with IoT data solutions for depression: A systematic review. *IEEE Journal of Biomedical and Health Informatics, 29*(3), 1919–1930.**

   This publication provides the evidence base for the thesis. It identifies the current capabilities and limitations of IoT-based systems for depression and emotional well-being and motivates the need for a more integrated, personalised, secure, and architecture-level approach [20].

2. **Zamani, S., Nguyen, M., & Sinha, R. (2025). Integrating environmental data for mental health monitoring: A data-driven IoT-based approach. *Applied Sciences, 15*(2), 912.**

   This publication examines environmental context as an underdeveloped dimension of IoT-based mental health monitoring. It contributes to the thesis by demonstrating how environmental data can be integrated into monitoring and interpretation processes [24].

3. **Zamani, S., Rostami, A., Nguyen, M., Sinha, R., & Madanian, S. (2025). A gamified AI-driven system for depression monitoring and management. *Applied Sciences, 15*(13), 7088.**

   This publication examines engagement and self-management through a gamified AI-IoT system. It contributes to the thesis by addressing the need for user interaction, adaptive feedback, and sustained participation in the context of depression monitoring and management [25].

4. **Zamani, S., Sinha, R., Madanian, S., & Nguyen, M. (2025). IoTMindCare: An integrative reference architecture for safe and personalized IoT-based depression management. *Sensors, 25*(22), 6994.**

   This publication forms the architectural core of the thesis. It integrates insights from the previous stages into *IoTMindCare*, a reference architecture that combines multimodal sensing, edge-cloud processing, personalised analytics, anomaly detection, feedback, safety considerations, and stakeholder-centred design [26].

5. **Zamani, S., Sinha, R., Madanian, S., & Nguyen, M. An expert-informed architecture for deployment-oriented IoT-based depression management. Submitted to Journal of Healthcare Informatics Research.**

   This manuscript provides the final refinement stage of the thesis. It incorporates technical and clinical expert perspectives to refine *IoTMindCare* with stronger attention to feasibility, calibration, alert governance, clinical workflow relevance, safety, privacy, interoperability, and future implementation considerations.

The relationship among the publications is shown in Table 1.1, which summarises their respective contributions, supporting evidence, and implementation boundaries. The first publication identifies the literature-derived gaps and design requirements (RQ1); the second and third explore contextual sensing and engagement-oriented design (RQ2); the fourth integrates these insights with the broader literature-derived requirements into IoTMindCare reference architecture (RQ3); and the fifth refines the architecture through technical and clinical expert feedback (RQ4). The table also distinguishes between literature synthesis, experimental and simulation-based evaluation, architectural design and analysis, and expert-informed refinement. These contributions collectively establish the foundation for IoTMindCare.

1.6 Organisation of the Thesis

This thesis is organised as a thesis by publication. Because the thesis includes published and manuscript-based chapters, prelude chapters are included to connect the publications into a unified doctoral argument. These preludes explain why each publication is included, how it relates to the central research question, how it connects to the preceding and following chapters, and how it contributes to the development and refinement of *IoTMindCare*.

Chapter 1 introduces the background of the thesis, the problem statement, the research questions, the methodological overview, the contributions, and the structure.

Chapter 2 provides a prelude to Chapter 3, which is a systematic review of IoT-based systems for depression and emotional well-being. This chapter identifies current capabilities and limitations in sensing, analytics, engagement, privacy, security, interoperability, and system integration, and establishes the need for a more coherent architecture-level response [20].

Chapter 4 provides a prelude to the environmental monitoring study presented in Chapter 5, which discusses a data-driven, IoT-based approach to environmental monitoring. This chapter examines temperature and humidity as contextual variables and demonstrates how environmental data can be incorporated into IoT-based mental health monitoring and interpretation [24].

Table 1.1 Research contributions, evidence, and implementation boundaries across the thesis.

Thesis stage	**Chapter/ Publication**	**Main contribution and output**	**Evidence and implementation boundaries**
Evidence base (RQ1)	Chapter 3 [20]	Identifies fragmentation and gaps in sensing, analytics, personalisation, safety, engagement, interoperability, and system integration. Provides literature-derived requirements for subsequent design stages.	Literature-based evidence.
Contextual sensing (RQ2)	Chapter 5 [24]	Investigates temperature and humidity data processing and classification as potential contextual inputs for IoT-based mental health monitoring.	Machine learning evaluation using environmental data and predefined stress-related proxy labels.
Engagement and self-management (RQ2)	Chapter 7 [25]	Develops a gamified research prototype and explores adaptive self-reporting, behavioural indicators, and feedback through simulation-based evaluation.	Prototype development and synthetic-data evaluation.
Architecture development (RQ3)	Chapter 9 [26]	Proposes the initial IoTMindCare reference architecture, integrating multimodal sensing, edge-cloud processing, personalisation, safety mechanisms, feedback, and stakeholder interfaces.	Architecture views, quality-attribute scenarios, and ATAM-informed trade-off analysis.
Expert-informed refinement (RQ4)	Chapter 11	Refines IoTMindCare using technical and clinical perspectives, introducing an explicit analytics and calibration layer and a hybrid edge-cloud processing view.	Questionnaire-based evaluation involving 22 experts; supports expert-informed refinement.

Chapter 6 provides a prelude to Chapter 7, which presents the gamified AI-driven system for depression monitoring and management. This chapter explores how micro-games, adaptive surveys, sensor data, personalised feedback, and AI analytics can support user engagement and self-awareness in depression monitoring and management [25].

Chapter 8 provides a prelude to the *IoTMindCare* reference architecture presented in Chapter 9, which serves as an integrative reference architecture for safe, personalised, and context-aware IoT-based depression monitoring and management. This chapter forms the architectural core of the thesis by organising multimodal sensing, edge-cloud processing, personalised analytics, anomaly detection, feedback mechanisms, safety considerations, and stakeholder roles into a coherent framework [26].

Chapter 10 provides a prelude to the expert-informed refinement study. Chapter 11 presents the expert-informed refinement of *IoTMindCare*. This chapter incorporates technical and clinical expert perspectives to strengthen the architecture in relation to feasibility, safety, calibration, alert governance, workflow relevance, interoperability, privacy, and future implementation considerations.

Chapter 12 provides an integrated discussion across the thesis. Rather than repeating the individual publications, it synthesises the findings around the major thesis-level themes: architecture-level integration, contextual interpretation, safety, personalisation, engagement, human oversight, alert governance, interoperability, and future adaptability.

Chapter 13 concludes the thesis by summarising the main contributions, acknowledging limitations, and outlining future research directions, including implementation, simulation, digital twin modelling, pilot studies, longitudinal evaluation, clinical validation, and further investigation of privacy, governance, and interoperability requirements.

The appendices provide supporting documentation for the human-participant component of the research. Appendix A contains the ethics approval documentation, Appendix B contains the participant information sheet, Appendix C contains the questionnaire used in the expert-informed evaluation, and Appendix D contains the feedback reports provided to participants following the questionnaire.

1.7 Chapter Conclusion

This chapter introduced the background, motivation, problem formulation, research questions, methodological overview, contributions, and thesis structure. It argued that IoT-based depression monitoring and management require more than isolated sensing, prediction, or prototype development. The central problem addressed by the thesis is the fragmentation of existing work and the lack of a coherent architecture-level foundation for integrating sensing, contextual interpretation, safety, personalisation, engagement, stakeholder interaction, interoperability, and governance.

In response to these gaps, this thesis develops and refines *IoTMindCare* as an integrative reference architecture rather than as a deployed or clinically validated intervention. The next chapter provides the prelude to the systematic review, establishing the evidence base and identifying the key gaps addressed throughout the thesis.

Chapter 2

Prelude to the Systematic Literature Review

2.1 Introduction

This chapter introduces the first major research stage of the thesis: a systematic literature review of IoT-based data solutions for emotional well-being and depression. The purpose of this prelude is to position the review within the overall thesis and explain why it is necessary before progressing to the later empirical, design-oriented, and architecture-focused chapters.

As discussed in Chapter 1, depression is a complex and contextual mental health condition. Supporting individuals with depression requires more than isolated symptom detection or occasional clinical assessment. It requires continuous, personalised, and context-aware approaches that can collect meaningful data, interpret changes over time, support user engagement, and address safety, privacy, and practical implementation concerns. IoT technologies offer significant opportunities in this area by enabling the collection of real-time data from wearable devices, mobile phones, smart-home sensors, and user-facing applications. However, the existing research remains fragmented across different technologies, data types, analytical methods, and intervention goals.

For this reason, the thesis begins with a systematic literature review. The review presented in Chapter 3 examines how IoT has been used to collect and analyse data for emotional well-being and depression. It focuses on data collection methods, data analysis techniques, emotion detection and monitoring, therapeutic approaches, depression prediction, depression detection, and depression management. In doing so, it establishes the current state of the field and identifies the gaps that motivate the remaining chapters of the thesis.

2.2 Role of the Review in the Thesis

The systematic literature review plays a foundational role in shaping the direction of the thesis. The main contributions of the literature review are:

1. Clarifying what kinds of IoT-based solutions already exist for emotional well-being and depression.

2. Identifying the types of data commonly used in these systems, including physiological, behavioural, self-reported, and environmental data.

3. Examining how analytical methods have been applied to interpret these data.

More importantly, the review reveals that many existing studies focus on specific components of the problem, such as emotion recognition, depression prediction, wearable sensing, mobile applications, or smart-home monitoring. While these contributions are valuable, they often remain disconnected from one another. Fewer studies provide an integrated approach that brings together sensing, analytics, personalisation, user engagement, safety mechanisms, privacy, interoperability, and practical depression management.

This gap is central to the motivation of this thesis. This research aims to develop a more integrated IoT-based approach for depression monitoring and management. Therefore, the systematic review provides the evidence base for the need for an architecture-level contribution.

2.3 Post-review Literature Update

The systematic literature review presented in Chapter 3 was last updated on 15 May 2024. Since then, further studies have been published on IoT-based mental health monitoring, passive smartphone and wearable sensing, multimodal data analysis, and AI-supported depression detection and monitoring. To ensure that the thesis remains current, a targeted post-review literature update was conducted for studies published after the original review cut-off date.

This update used the same conceptual scope as the original review, focusing on IoT, sensor-based behavioural and physiological monitoring, wearable and smartphone sensing, contextual or environmental sensing, data collection, data analysis, depression, and emotional well-being. Broad literature searches, particularly in Google Scholar, returned more than 1000 records after the original review cut-off date. However, many of these records were outside the specific scope of this thesis, duplicated earlier themes, focused on general digital

mental health without IoT-based sensing, or addressed mental health conditions without a clear depression or emotional well-being focus. After narrowing the results to studies directly relevant to IoT-based or sensor-supported depression and emotional well-being monitoring, 13 recent studies were identified as highly relevant. Because this update was conducted as a targeted post-publication scan rather than a full database-native PRISMA rerun, the results are used here to confirm the currency of the thesis direction rather than to replace the systematic review in Chapter 3.

Table 2.1 summarises the main directions identified across these 13 directly relevant post-cut-off studies and shows how they relate to the gaps identified in the original systematic review.

Table 2.1 Summary of post-review literature developments identified from 13 directly relevant studies published after 15 May 2024

Recent development	**Main observation in recent studies**	**Relevance to this thesis**
Smartphone and wearable sensing	Recent studies continue to use activity, step count, heart rate, sleep, phone usage, location, and speech or language data for depression-related monitoring [27, 15, 28].	Confirms the continued importance of sensor-based behavioural and physiological monitoring.
Multimodal and longitudinal data	Newer studies increasingly combine smartphone, wearable, physiological, behavioural, and self-reported data over time [29–33].	Supports the thesis argument that depression monitoring requires integrated and longitudinal data rather than isolated measures.
AI-supported analysis	Machine learning and AI remain dominant for depression screening, symptom monitoring, and prediction [34, 35, 27].	Reinforces the need to connect data analysis with safe, explainable, and clinically informed system design.
Contextual and environmental sensing	Environmental variables such as temperature, humidity, weather, and ambient conditions are beginning to appear, but remain less developed than physiological and behavioural data [24].	Supports the contribution of the contextual sensing chapter and the need to integrate environmental context into IoT-based depression monitoring.
Management and intervention support	Most studies still focus on detection, prediction, or monitoring rather than active long-term depression management [25, 26].	Justifies the thesis focus on user engagement, personalisation, adaptive feedback, and architecture-level integration.
Privacy, interoperability, and deployment	Recent studies continue to report concerns about sensitive data, consent, anonymisation, validation, interoperability, and clinical deployment [15, 31].	Strengthens the need for an integrated reference architecture that addresses technical, ethical, and practical requirements.

The newer literature largely reinforces the original systematic review's findings. The targeted post-review update identified 13 recent studies that were directly relevant to the scope of this thesis. These studies indicate continuing research interest in IoT-based and sensor-supported approaches to depression and emotional well-being, while remaining focused enough to be discussed as a post-review update rather than a second systematic review. As shown in Table 2.1, recent studies continue to show strong interest in smartphone and wearable sensing for monitoring depression and emotional well-being. Commonly used data sources include accelerometer-derived activity, step count, heart rate, sleep, app usage, call behaviour, location or mobility features, ecological momentary assessment, and speech or language data. Recent reviews also confirm that AI and machine learning remain dominant analytical approaches, particularly for depression screening, symptom monitoring, and sensor-based behavioural and physiological monitoring [27, 15, 28, 34, 35].

A noticeable development since the original review is the growing emphasis on longitudinal and multimodal sensing. Newer studies increasingly combine smartphone, wearable, behavioural, physiological, and self-reported data to track depression severity, identify markers of treatment response, or model individual variation over time [29–33]. This strengthens the argument that passive and multimodal sensing can provide useful signals for depression-related monitoring. However, the evidence also remains methodologically uneven. Many studies still report small or moderate sample sizes, short monitoring periods, limited external validation, heterogeneous feature sets, and limited clinical deployment.

Importantly, the update confirms that several gaps identified in the original review remain open:

1. While multimodal sensing is becoming more common, there is still limited agreement on clinically generalisable features, validation protocols, and deployment standards.

2. Environmental and contextual sensing remains comparatively underdeveloped. Although recent work has begun to examine variables such as temperature, humidity, weather, and ambient conditions, these factors are still not routinely integrated into IoT-based depression monitoring systems [24].

3. The literature continues to focus more heavily on detection, prediction, and monitoring than on active long-term depression management. Only a small number of recent studies address personalised intervention, adaptive feedback, user engagement, or architecture-level integration [25, 26].

Privacy, security, interoperability, and design-oriented implementation considerations also remain important limitations in the recent literature. Studies using richer data streams, such

as speech, location, and behavioural traces, highlight the continuing difficulty of balancing data utility with re-identification risk, user consent, anonymisation, and responsible data sharing [15, 31]. These concerns are particularly important in depression-related systems because the data collected are sensitive and may influence clinical decision-making, self-management, or caregiver involvement.

Overall, the post-review literature update supports the continued relevance of the original systematic review and strengthens the motivation for this thesis. The field has progressed toward more longitudinal, multimodal, and AI-supported sensing, but it remains fragmented across devices, datasets, algorithms, and intervention strategies. Environmental/contextual data, active depression management, privacy-aware implementation, interoperability, and architecture-level integration remain insufficiently developed. These findings further justify the need for the subsequent chapters of this thesis, which examine contextual sensing, user engagement, and the development of *IoTMindCare* as an integrated reference architecture for safe and personalised IoT-based depression monitoring and management.

2.4 Connection to the Following Chapters

The findings of the systematic review directly inform the later stages of the thesis. The review shows that environmental and contextual data are relevant to emotional well-being and depression, but they are not yet sufficiently integrated into IoT-based mental health systems. This motivates the following research stage, which examines contextual sensing through environmental variables such as temperature and humidity.

The review also highlights the need for more active and engaging systems. Many IoT-based approaches rely heavily on passive sensing or detection, but depression management also requires user interaction, feedback, and support for self-management. This motivates the later study on a gamified AI-driven system for depression monitoring and management.

Finally, the review identifies broader challenges related to privacy, security, interoperability, personalisation, data quality, and practical implementation. These issues inform the development of the *IoTMindCare* architecture, which aims to integrate the technical, contextual, ethical, and practical requirements of IoT-based depression management into a coherent reference architecture. In this way, the systematic literature review provides the foundation for the remaining chapters of the thesis, which progressively address contextual sensing, user engagement, architecture-level integration, and expert-informed refinement.

Chapter 3

Enhancing Emotional Well-Being With IoT Data Solutions for Depression: A Systematic Review

3.1 Abstract

Effectively caring for adults with depression is challenging. While technology offers potential improvements in emotional well-being through better monitoring, standardised methods to gather and analyse relevant data are highly fragmented. This Systematic Literature Review (SLR) explores using Internet of Things (IoT) based data collection and analysis to enhance emotional well-being and manage depression effectively. Following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) framework, we report in-depth findings from 42 studies, which were selected from an initial set of 559 published works. We find that current literature extensively covers important topics like IoT for detecting, analysing, and monitoring emotions, therapeutic interventions for emotional well-being, and predicting, detecting, and managing depression. IoT-based data collection and analysis solutions predominantly employ sensors and AI, respectively. The literature review identifies a gap in prioritising active systems that engage users, highlighting the need to address key aspects such as privacy and security.

Publication Note

This chapter is based on the published article:

Zamani, S., Sinha, R., Nguyen, M., & Madanian, S. (2025). Enhancing emotional well-being with IoT data solutions for depression: A systematic review. IEEE Journal of Biomedical and Health Informatics, 29(3), 1919-1930.

The published article is available online via DOI: 10.1109/JBHI.2024.3501254.

The remainder of this chapter is based on the published article, with minor formatting adjustments made to align it with the structure and argument of this thesis.

3.2 Introduction

Depressive disorder or *depression* is a widespread mental health issue that arises from a multifaceted interplay of social, psychological, and biological elements [4]. Recently, due to increasing stress levels, lifestyle changes and disruptions like COVID-19, the incidence of depression has increased dramatically [5]. Monitoring and supporting adults with depression requires finding scalable, efficient and effective solutions.

Recent advances in technology, like the proliferation of the IoT, have the potential to address mental health issues, including depression [36]. IoT provides several benefits over traditional approaches like clinic-based care or over-the-phone remote care. These benefits include scalability, cost-effectiveness, and personalisation. IoT-based solutions can capture relevant patient information in real-time with minimal effort, improving both the quality and quantity of data to support patients.

Managing depression requires collecting and analysing relevant data, which may differ from other mental health disorders. This article specifically focuses on IoT solutions for depression and emotional well-being. Despite significant developments in the use of IoT for supporting emotional well-being and depression, there is no current secondary study that comprehensively covers the area of data collection and analysis for supporting emotional well-being and depression. Table 3.1 lists past surveys relevant to this subject, highlighting this gap in the literature, which we aim to address through a systematic literature review (SLR). While concerns like privacy and security are critical, our SLR is scoped to focus only on data collection and analysis for emotion detection, monitoring and analysis techniques, therapeutic approaches, and methods for predicting, detecting and managing depression.

We explore our primary research question: *"In what ways is IoT being used to collect and analyse data to improve and monitor the emotional well-being of adults with depression?"* This SLR offers valuable insight for both technology developers and healthcare providers interested in leveraging existing IoT solutions in this application area.

We follow the SLR methodology [49] described in Section 11.4. Sections 3.4 and 3.5 present the data collection and analysis methods we reviewed, respectively. In Section 3.6,

Table 3.1 Comparison of the coverage of our research with existing surveys

Source	Data Collection			Data Analysis			Emotional Well-being			Depression		Scope of Review
	Sensor-based	User-centric	Datasets	AI-based	Big data	Statistics	Detection	Monitor/ Analyse	Therapy	Prediction/ Detection	Management	
[37]	✓	✓	✓	✓			✓	✓	✓	✓		Comparing traditional methods with emerging technology to improve mental health.
[38]	✓			✓	✓	✓		✓	✓	✓		Autonomous mental health monitoring, with real-time analysis, and data integration.
[39]				✓						✓		Textual emotion detection, focusing on NLP, ML, and ethical considerations.
[40]	✓		✓	✓				✓	✓	✓		Physiological emotion recognition using real-life data from wearables.
[41]								✓	✓	✓		IoT-based telemedicine for enhancing disease prevention and health promotion.
[42]	✓	✓						✓	✓	✓		Sensor usage in e-health and m-health for diagnosing mental health issues.
[43]			✓	✓				✓		✓		Real-time mental health detection and analytics using IoMT and social media data.
[44]	✓	✓		✓				✓	✓	✓		Smartphone-based technologies for mental health, focusing on COVID-19 impacts.
[45]	✓	✓						✓		✓		Intelligent positive computing with IoT devices, for improving well-being.
[46]	✓				✓			✓	✓			IoT technology for patient monitoring and data analysis, with AI and cloud computing.
[47]	✓			✓				✓		✓		Routing protocols in IoT-based wireless sensor networks for healthcare.
[48]	✓			✓				✓	✓	✓		Using multimodal data fusion in predictive, preventive, and personalised healthcare.
Our review	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	Data collection and analysis methods for emotional well-being and depression.

we discuss approaches to promoting emotional well-being, including strategies, interventions, and frameworks. Section 3.7 focuses specifically on depression, exploring interventions and methodologies to address this complex mental health condition. Section 11.6 offers a comprehensive discussion synthesising key findings, identifying patterns, and drawing connections. Section 11.7 concludes with a summary of our main insights, implications, and future works.

3.3 Systematic Literature Review Methodology

Our SLR adheres to the PRISMA guidelines [50], encompassing initial planning, systematic search and retrieval of the literature, rigorous screening and selection, meticulous data extraction, detailed quality assessment, comprehensive data synthesis and analysis, followed by report writing.

3.3.1 Planning, Search, and Retrieval

We combined the keywords relating to our primary research question into the following search string:

> *("IoT" OR "Internet of Things") AND ("emotional well-being" OR "emotional wellbeing" OR "emotional well being") AND "depression" AND ("data collection" OR "data analysis") AND ("effectiveness" OR "effective")*

This search string was applied to both titles and abstracts of the papers available in the following databases: IEEEXplore, Science Direct, Scopus, Springer, AIS eLibrary, ACM Digital Library, and EBSCO. The annual distribution of our search results is illustrated in Figure 3.1. A significant rise in the number of studies since 2015 informed our decision to limit our search to literature published between 2015 and 2024. However, we also considered seminal works outside this date range, where required, through snowballing from our initial search results.

3.3.2 Screening and Selection

To organise and manage search records and references efficiently, we employed Excel and Mendeley. The studies were last organised on 2024-05-15. A set of 559 studies were initially exported as text files and Microsoft Excel spreadsheets, sorted by titles, authors, years, sources, and keywords.

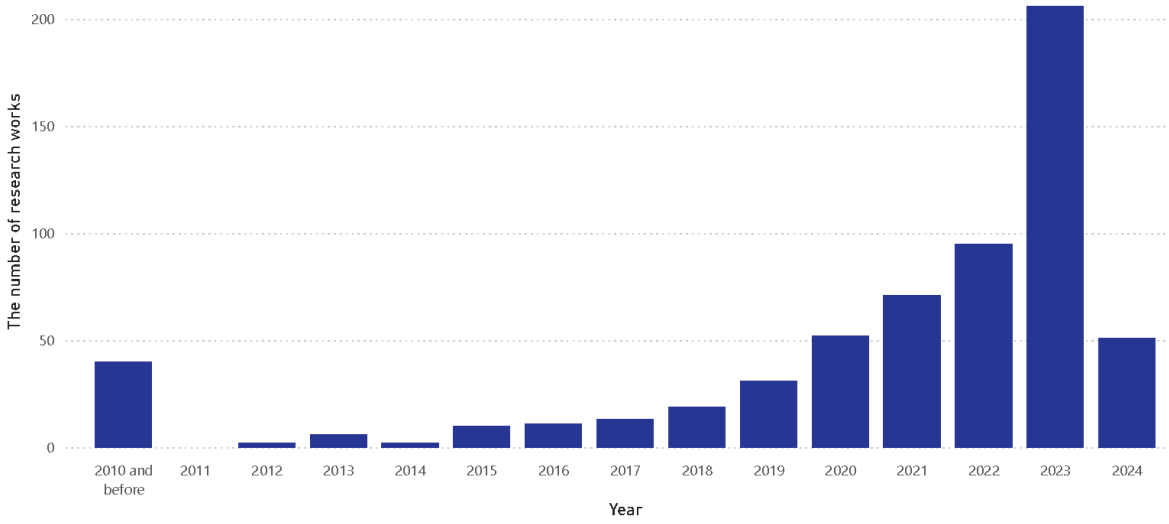


Fig. 3.1 Publications per year for the provided search string

The initial screening involved reviewing titles and keywords to determine their relevance to our study's scope, reducing the number to 169 potentially relevant studies. The subsequent stage entailed a more detailed review of the studies' abstracts, categorising them according to the title, problem statement, solution, study aims, and reasons for inclusion or exclusion. The status field was labelled as YES (accepted for further review), MAYBE (possibly accepted for further study), and NO (rejected for further study based on our criteria). This process further refined the selection to 90 studies. Table 3.2 provides our overall inclusion and exclusion criteria.

Table 3.2 Inclusion and exclusion criteria

Criteria	**Inclusion**
Date	Publications between 2015-2024
Language	Studies written in English
Topic	Studies covered IoT, emotional well-being, and depression
Accessibility	Studies with full-text availability
Peer Review	Peer-reviewed articles

In the final analysis, we excluded any work not aligning with our research questions by meticulously reading each study's introduction and conclusion. This rigorous process led to a concentrated group of 42 research articles addressing our research questions directly. This overall process followed to identify the final set of studies is illustrated in Fig. 3.2.

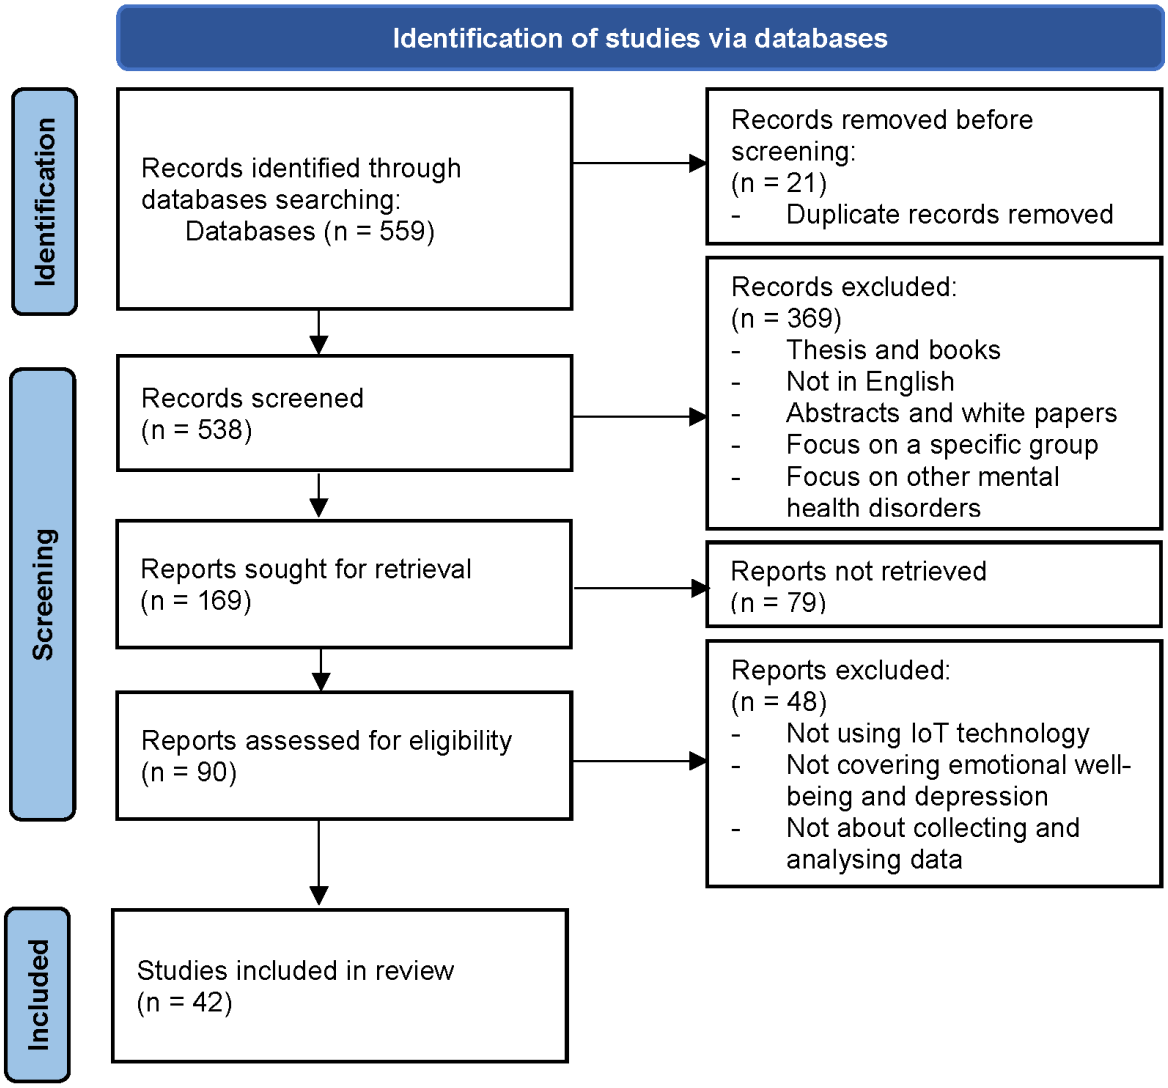


Fig. 3.2 Systematic Literature Review (SLR) process

3.3.3 Data Extraction and Analysis

The critical aspects of our primary research question were explored through the following sub-questions to systematically analyse the selected works and present our findings in subsequent sections: *(1) What data collection methods are being used?; (2) What data analysis methods are being used?; (3) What emotion detection, monitoring, and/or analysis do the studies cover?; (4) What therapeutic approaches for emotional well-being are being used, if any?; (5) What techniques for predicting and detecting depression are being used?; (6) What strategies for managing depression are presented?*

3.4 Data Collection Methods

Patient data pertaining to depression and emotional well-being encompasses a broad spectrum of physical and mental measurements, such as personal details, medical history, and brain imaging, as noted in the literature [51]. The literature identifies three data collection categories: *sensor-based*, *user-centred*, and *dataset-based*. Figure 3.3 shows the distribution of studies across these categories. Wearable technologies and IoT are the most popular, followed by using questionnaires and mobile phone sensors.

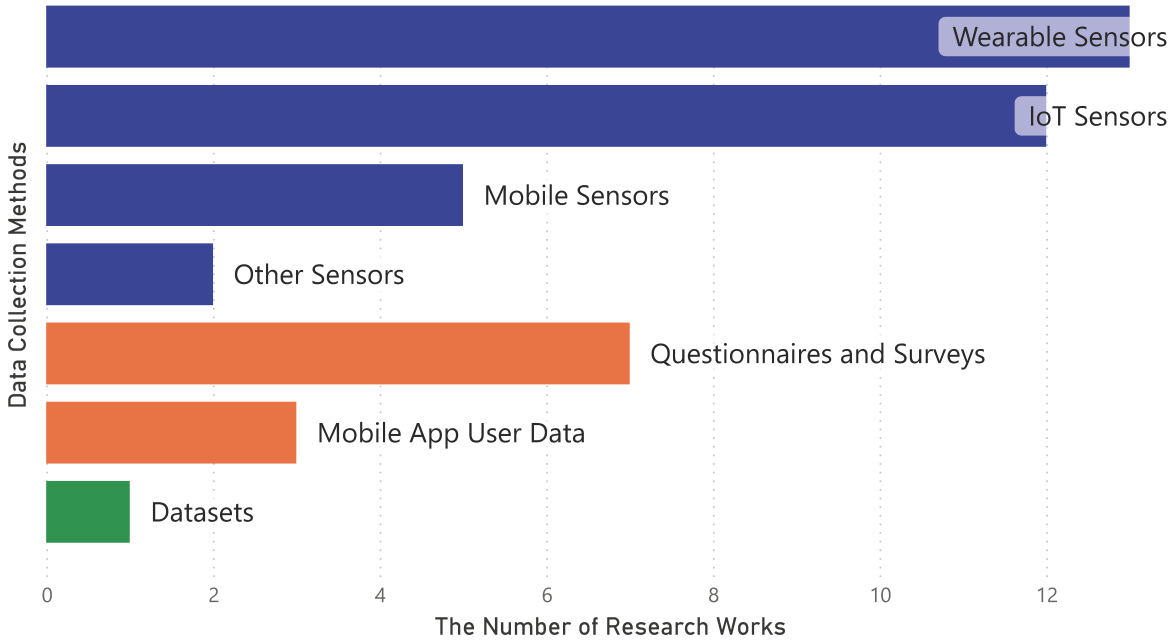


Fig. 3.3 Number of data collection methods (blue: sensor-based methods, orange: user-centric methods, green: datasets)

3.4.1 Sensor-based Data Collection

Sensor-based data collection offers several advantages. Sensors can continuously gather data non-intrusively [52] from mobile applications, multi-modal physical devices, as well as text, speech, images and videos [37]. Cost, size and intrusiveness are key factors in selecting sensors [47], and these often pose barriers to making consistent care available to larger community groups. Wearable, IoT, and mobile phone sensors are the most popular types identified in the literature.

Wearables

Wearables, like smartwatches and bracelets, can gather accurate and real-time quantitative data that can improve physical health and promote positive habits [40]. Direct access to personal analytics and data can be instrumental in self-management and facilitating preventative care [53]. Wearables can effectively capture data on physical activity, sleep cycles, and other health-related parameters [54]. Table 3.3 catalogues studies that utilise wearables for data collection.

Table 3.3 Data collection via wearables, categorised by type of data, wearable type, and purpose of data collection

Data	Source	Wearable Type	Purpose
Physiological	[55]	Wearable gadgets	Discern emotions
	[53]	Commercial wearables	Remote health monitoring
	[56]	Wearable sensors	Monitoring depression level
	[57]	Smartwatch	Consistently monitor users
	[58]	Wearable sensors	Detect anxiety by emotions
	[59]	Wearable devices	Monitoring indicators
Behavioural	[60]	multisensor wearable	Unobtrusive data collection
	[61]	IoT-based wearable	Monitor the suicidal ideation
Activity	[62]	Noninvasive IoT	Remote data monitoring
	[63]	Wrist and chest-worn	Mental stress state detection

Smart watches are the most used wearable in the literature and can capture heart rate, physical exertion, and sleep patterns [55, 53, 59]. More bespoke wearables like skin sweat sensors provide physiological data related to emotional states [42]; multi-sensor-integrated wearables offer even more complex data collection [60] but cost more and less ubiquitous.

IoT Sensors

IoT sensors integrated within smart homes facilitate tracking daily routines, behaviours, and contextual factors relevant to mental health [46]. Motion sensors can monitor activity levels, potentially indicating changes such as a decline in motivation or social interaction, symptoms often associated with depression [64]. Table 3.4 lists the studies that use IoT sensors.

The Internet of Medical Things (IoMT) extends IoT and emphasises the critical role of information exchange [72]. IoMT sensors can enhance healthcare through a high degree of customisation in patient care [73]. Chen et al. [65] utilised IoMT sensors for collecting text, audio, and visual inputs to detect depression. IoMT helps professionals better address depression and improve emotional well-being by analysing physiological indicators. For example, the activation of the Sympathetic Nervous System during stress increases heart rate,

Table 3.4 Data collection by IoT sensors, categorised by type of data, sensor type, and purpose of data collection

Data	Source	IoT Sensor	Purpose
Images	[65]	IoMT sensors	Depression detection
	[66]	Smart home sensors	Analysing facial expressions
Activity	[67]	PIR sensors	Track and monitor elderlies' daily behaviour
Physiological	[68]	Embedded sensors	Identify sociability patterns
	[69]	EmoGo IoT device	Monitor and categorise emotions
	[70]	WBAN sensors	Mental health monitoring
	[71]	IoT sensors	Integrate connected sensors

respiration rate, and sweat secretion, impacting various physiological systems to respond to challenges [40]. A decrease in Heart Rate Variability indicates a disruption in the Autonomic Nervous System balance, which is associated with the inability to manage stressors linked to depression [74]. Such complex physiological phenomena can be captured effectively by IoMT sensors [69–71].

Environmental factors can also significantly affect mental health. Indoor conditions can be tracked using sensors embedded in objects, like smart beds and chairs [38]. Several works integrate IoT-based home and environmental data to enhance depression management [67, 69]. Temperature is a significant environmental variable. Some studies show that extremely hot or cold temperatures may adversely impact mental health [75].

Mobile Phone Sensors

Smartphones are widely available and equipped with an array of sensors capable of collecting real-time data for personalised, responsive and high-quality healthcare. Accelerometer data from phones can be used to identify critical information like physical activity patterns, sleep quality, falls, and voice characteristics [59]. Other phone sensors like GPS and motion sensors have also been used to monitor depression [44].

Other Sensors

Some specific data can only be collected by bespoke sensors, such as colour sensors for monitoring moods and emotional states [42], and robotic devices within Body Area Networks to capture text, audio, video and images to detect depression [76].

3.4.2 User-generated and User-centric Data Collection

User-generated data collection is one of the most widely used methods for predicting, detecting, and managing depression and emotional well-being concerns. Data collected directly from users can offer deep insights into their emotional states. Popular mechanisms include questionnaires and interviews, and non-intrusive means such as observing users' social media activities and screen time or phone usage patterns. However, human factors like mood fluctuations and biases can affect the reliability of user-generated data.

Questionnaires and Surveys

Questionnaires and surveys are a flexible instrument to collect a variety of data. They are extremely popular in many depression-related healthcare systems, as shown in Table 3.5. With increasing digitisation, traditional paper-based questionnaires are being replaced by their digital counterparts [44, 45].

Table 3.5 Data collection by questionnaires and surveys, categorised by question types or nature, and their aim of collecting data

Source	Question Types	Purpose
[62]	Pre, mid, and post assessment	Collecting relationship measures
[67]	Included GDS	Assess depressive symptoms
[77]	Linked data points	Construct predictive model
[60]	Self-assessment	Evaluate psychological and mental status
[78]	Smartphone-based	Smartphone and emotion relationship
[71]	VR questionnaires	Comparing responses with outcomes
[57]	Standard questionnaires	Utilise tagged data from smartwatches

Mobile Applications

Mobile applications have emerged as a popular choice for digitising traditional questionnaires, providing convenience, accessibility and non-intrusiveness. Research indicates an inverse relationship between social activity and depression levels [42]. As depression and anxiety often result in social isolation, tracking social activity becomes invaluable [42]. Analysing text sources such as social media posts can also offer valuable data related to mental well-being [37]. Table 3.6 shows works that use mobile applications for data collection.

Table 3.6 Data collection by mobile applications, categorised by type of data and method of collection

Source	Data Type	Method
[62]	Daily surveys	Mobile applications to monitor mental health
[55]	Digital interactions	Observing online behaviours using ML
[59]	Social data	Analyse social networks, for healthcare data

3.4.3 Datasets

Utilising diverse datasets can enhance the reliability, robustness and effectiveness of solutions for depression. Research contains a mix of lab-based synthetic datasets as well as those from real-life contexts [40]. Popular datasets containing physiological data related to emotional states include DAPPER [79], AMIGOS [80], and CASE [81]. Moreira et al. [82] focused on predicting postpartum depression (PPD) and utilised genuine patient datasets from the Maternity School Assis Chateaubriand. This approach underscores the value of integrating real-life patient data into their analytical framework. Relying entirely on datasets can be insufficient, as they may not fully reflect real-world conditions in specific applications. It is recommended that datasets be complemented with other data collection methods.

IoMT datasets contain data obtained from monitoring tools, wearables, and other sensors, while social media datasets comprise user-generated content such as posts, comments, and videos [43]. Solutions that use multiple datasets may benefit from improved efficiency, accuracy and adaptability.

3.5 Data Analysis Methods

Raw data contains low-level information [67] that can be analysed to provide richer insights. Data collection and analysis complement each other in developing effective systems. Data analysis can lead to solutions for proactive healthcare management, early detection of diseases, and personalised treatment approaches [47]. Data analysis methods can be categorised under *AI-based, Big Data, statistical,* and *advanced data analysis*. Figure 3.4 shows the number of studies in each category; AI-based methods are by far the most popular.

3.5.1 AI-based Methods

AI and Machine Learning (ML) algorithms promise the ability to accurately analyse and classify large volumes of complex data. Figure 3.5 illustrates the frequency of emotional

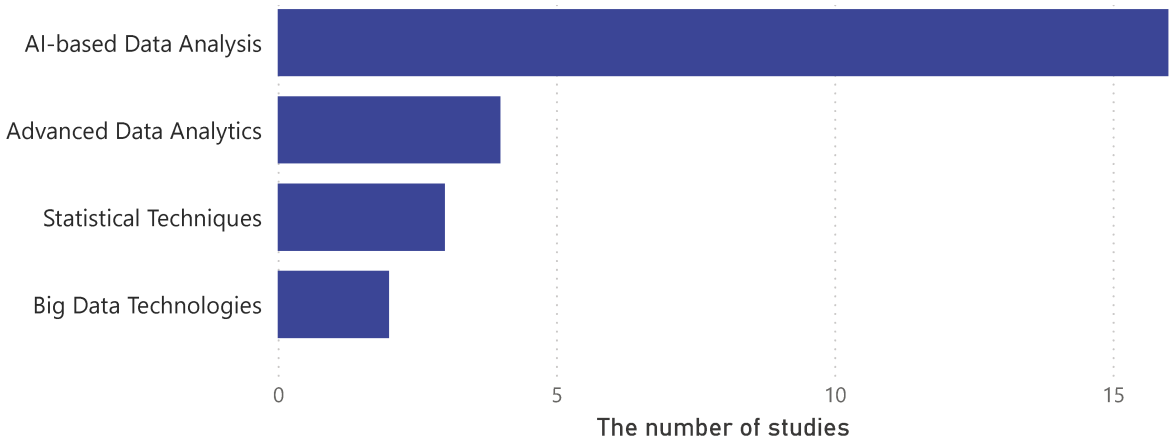


Fig. 3.4 Data analysis methods in emotional well-being and depression healthcare systems

well-being solutions over time, incorporating AI-based analysis. Table 3.7 summarises AI-based data analysis methods.

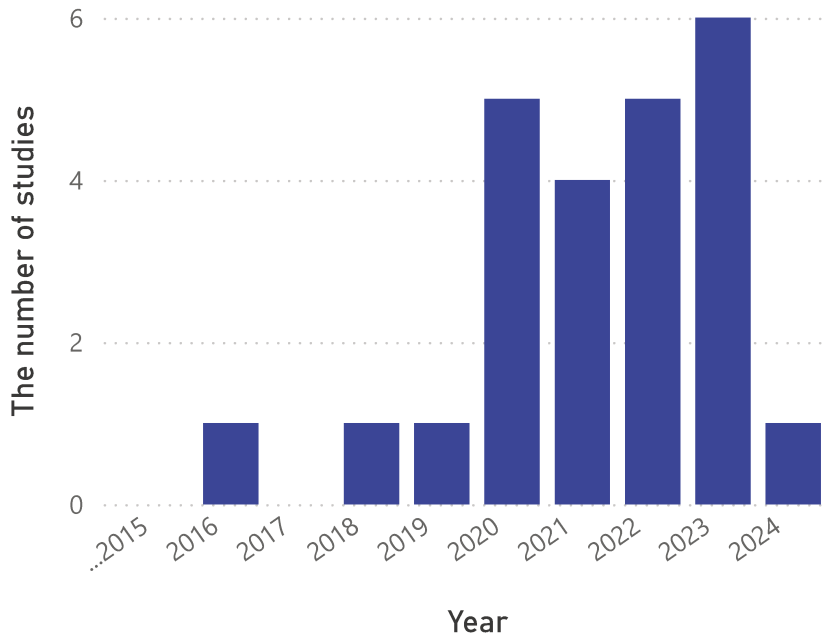


Fig. 3.5 Trend of using AI for analysing data in emotional well-being and depression healthcare solutions including surveys

Natural Language Processing (NLP) plays a vital role in this area. By identifying linguistic patterns and assessing emotional states from communications, NLP can tailor interventions to support individuals better. An example is Textual Emotion Detection, which can automatically analyse and identify predominant emotions in text [39].

Table 3.7 AI-based solutions categorised by type of algorithm and purpose of analysis

Source	AI Algorithm Type	Purpose
[77]	Clustering	Develop classifier models
[69]	Decision-making	Training emotional states
[83]	CNN and ML	Classifies and analyse emotions
[55]	Bi-LSTM	Uncover depression and stress patterns
[82]	Decision trees	Predict PPD among pregnants
[84]	KNN and NLP	Analyse and interpret mental health data
[85]	CNN	Accurate emotion recognition
[68]	FIS, and ML	Detect alterations in social behaviour
[76]	Robotic-based	Detecting depression rates
[63]	DNN	Classify and detect mental stress state
[86]	DL	Process and classify health data
[71]	ML	Process, format, and classify data
[57]	ML	Predict users' wellness
[66]	ML	Identify and classify emotions
[61]	LR	Classify event data
[58]	Bayesian	Emotions recognition to detect anxiety

3.5.2 Big Data

Big Data technologies can perform analyses over large volumes of complex data, such as medical records and sensor readings, leading to more accurate diagnoses, tailored treatment plans, and timely interventions.

Analysing Big Data can pose challenges, requiring preprocessing techniques for data cleaning, normalisation, and feature extraction to enhance data quality and model performance [87]. These steps are critical in reducing noise and improving the accuracy and reliability of solutions. Ali et al. [59] employ several data analysis techniques to extract valuable features, simplify data, and categorise patient health conditions. Similarly, Rajawat et al. [76] explore the application of big data analytics to enhance the system's awareness of the patient and the broader contextual conditions for providing personalised care.

3.5.3 Statistical Techniques

Statistical methods assist in recognising patterns, correlations, and trends over time. They could be used to pinpoint risk factors, evaluate treatment efficacy, and forecast future outcomes. Techniques like Principal Component Analysis and multidimensional scaling are used for dimensionality reduction for efficient analysis [38]. Ou et al. [67] process sensor activations to monitor the daily movements of elderly individuals, identifying "transition

timestamps" to mark shifts between sensors. By analysing duration in different places, they extracted activity patterns indicative of depression.

Descriptive statistics, multilevel regression models, and statistical tests can explore the link between smartphone use, emotion regulation and mood changes. Shi et al. [78] used these techniques and thematic analysis to identify patterns and themes, providing a comprehensive explanation of participants' smartphone-based emotion regulation behaviours. Mano et al. [66] employed statistical analysis to validate findings and assess the effectiveness of various classification algorithms.

3.5.4 Advanced Data Analysis Methods

In this section, some more sophisticated methods for analysing data are discussed. In depression detection systems, fusion techniques integrate diverse data sources and enhance the accuracy of identifying depressive symptoms. As an example, Chen et al. [65] introduce the IIFDD (Intra and Inter-modal Fusion Model for Depression Detection) framework to detect depression. This framework combines multimodal feature extraction with intra-modal and inter-modal fusion techniques. Yang et al. [60] analyse the correlation between data collected from wearable devices and social behaviour, employing threshold selection, correlation analysis, long-term analysis, and classification methods. They use a feature fusion and classification algorithm to analyse wearable device data, focusing on preprocessing and calculating features for correlation and classification.

Employing cognitive tests alongside behavioural analysis can provide a comprehensive approach to monitoring mental health and understanding cognitive functioning in individuals. Yunusova et al. [62] utilise multilevel modelling, time series analysis, dynamic multilevel modelling, exploratory factor analysis, and data imputation. Additionally, multilevel modelling analyses associations within intensive longitudinal data, accommodating multiple evaluations per individual. Time series analysis and dynamic multilevel modelling techniques are commonly applied to examine temporal associations and trajectories between constructs of interest.

In the context of social media datasets, Gupta et al. [43] discuss the extraction and analysis of psycholinguistic features from user-generated content to discern behavioural patterns and emotional states. Moura et al.'s solution [68] achieves incremental learning of context-sensitive sociability patterns by combining Frequent Pattern Mining (FPM) and Complex Event Processing (CEP). FPM, a computational technique, aims to uncover frequent patterns in various data collection types, including relational and non-relational databases, text files, and data streams.

3.6 IoT-based Emotional Well-being

Emotional well-being strategies can ease depression symptoms and prevent future episodes. Solutions for enhancing emotional well-being can be organised into three distinct categories: *1) IoT-based, 2) IoT integrated with AI (AI-based),* and *3) User-centric* methodologies. Figure 3.6 shows the corresponding categorisation of relevant studies from our review. AI-integrated IoT methods are commonly used for emotion detection, monitoring, and analysis. These methods offer the ability to detect subtle emotional patterns and continuous personalised feedback. Key challenges include privacy issues, possible algorithmic biases, and dependence on high-quality training data, which may not always be fully representative. IoT-based solutions can collect real-time, continuous data from various sources and provide a more comprehensive view of a patient's emotional state. However, these solutions often require significant infrastructure investment and feature data security and device interoperability concerns. Finally, IoT integrated with user-centric approaches can enhance user experience by providing personalised interventions based on direct feedback. Challenges include subjective and/or inconsistent user inputs and variability in user interactions over time.

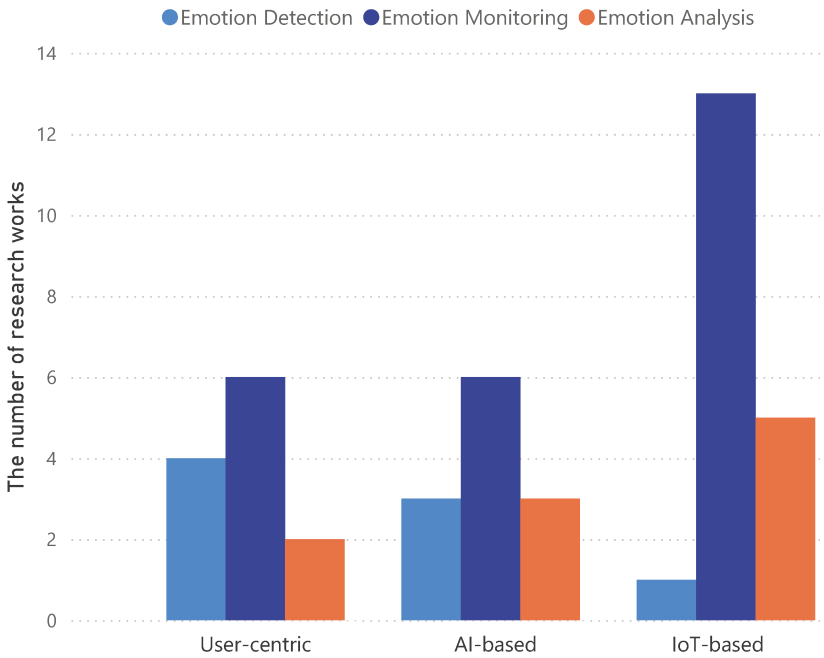


Fig. 3.6 Emotional well-being solutions based on types and methods

3.6.1 Emotion Detection

Emotion detection involves identifying and recognising emotions at a specific point in time through cues like facial expressions, tone of voice, body language, and physiological signals such as heart rate or skin conductance. Emotion detection serves as a key aspect in supporting well-being [37].

Several studies use ML to detect emotions. Gyrard et al. [88] use facial recognition and speech recognition, used by robots to understand human emotions. They use IoT to collect physiological and emotional data for analysis and interpretation. Mano et al. [66] utilise ML algorithms for monitoring patients and recognising emotions through facial expressions.

The M-SID framework [61] uses physiological signals to discern patterns of suicidal ideation and real-time detection of emotions. Mahdi et al. [69] delve into facial recognition, ECG, EEG, BPM, and temperature signals for emotion detection. The i-RISE framework [58] integrates eye-event tracking, heart rate monitoring, daily activity tracking, and ambient sensing to identify and interpret individuals' emotional states.

3.6.2 Emotion Analysis and Monitoring

Emotional analysis improves the accuracy and timeliness of detecting depression through contextual analysis, tracking patterns of emotions over time, and observing fluctuations in emotional well-being [38]. A categorisation of research works on emotion monitoring and analysis is shown in Table 3.8.

IoT-based Emotion Monitoring and Analysis

Gyrard et al. [88] classify and examine projects that improve emotional well-being using IoT devices and ontologies. They highlight the importance of including various IoT sensors to monitor physical and mental states. The research also underscores the role of IoT technology in showcasing knowledge of APIs.

MEmoR [85] is a multimodal ER model integrated with IoT devices and wearables to monitor users' emotional states using visual and physiological cues from cameras, microphones, smartwatches, and smartphones. The EmoGo system [69] tracks emotions using wearable data and categorises emotions before transmitting them to a server for further analysis.

Several works carry out emotion monitoring and analysis in smart home settings. Yunusova et al. [62] investigate the interconnections between daily variables and long-term mental health and well-being. They employ data reduction strategies and multilevel modelling to

Table 3.8 Methods for analysing and monitoring emotions

Method	Research Work	Analysis	Monitoring
IoT-based Methods	Kumar et al. [85]	✓	✓
	Mahdi et al. [69]		✓
	Gyrard et al. [88]		✓
	Kumar et al. [55]	✓	✓
	Yunusova et al. [62]		✓
	Ou et al. [67]		✓
	Yang et al. [60]		✓
	Jose et al. [70]		✓
	Gupta et al. [56]	✓	✓
	Chen et al. [65]	✓	
	Rajawat et al. [76]		✓
	Mano et al. [66]		✓
	Sundaravadivel et al. [58]	✓	✓
AI-based Methods	Moreira et al. [82]	✓	
	Ali et al. [59]		✓
	Sinha et al. [83]		✓
	Robinson et al. [53]		✓
	Hardy et al. [71]	✓	✓
User-centric Methods	Mele et al. [89]		✓
	Shi et al. [78]	✓	✓

examine the relationships between daily emotions, activities, and sleep patterns for individuals. Ou et al. [67] utilise IoT-enabled smart home devices to monitor residents' behaviours continuously. Mano et al. [66] track facial expressions over time to observe emotions through IoT sensors implemented in smart homes.

Some studies consider wearable devices as components of a larger IoT network. Yang et al. [60] read data from wearables and analyse it to monitor emotional well-being.

Some research works use bespoke IoT sensors. Gupta et al. [56] incorporate GSR and finger-measurable heartbeat sensors for monitor anxiety and depression levels. Jose et al. [70] explore integrating IoT technology with WBAN and biosensors for remote health monitoring purposes, enabling the tracking of patients' vital signals to support recovery. Kumar et al.'s study [55] suggests using data on emotional dynamics to design an IoT architecture model that integrates therapeutic approaches for emotional problem prediction and treatment. Sinha et al. [83] leverage facial recognition and emotion analysis to monitor students' emotional well-being in real-time.

The i-RISE framework [58] incorporates elements for eliciting emotions through a semi-immersive environment. Rajawat et al. [76] process data through fog and cloud computing

platforms and send automatic alerts to medical professionals to reduce the risk for older adults in critical situations.

AI-based Emotion Monitoring and Analysing

Some IoT solutions employ ML to assess collected data and identify individual mood trends [47]. Moreira et al. [82] investigate using ML techniques for pattern recognition to anticipate post-partum depression by analysing clinical and physiological data. Ali et al. [59] extract and filter patient messages and posts from social networks using supervised and unsupervised ML to classify the sentiment polarity of text. Hardy et al. [71] propose a method to track emotions via virtual reality and physiological signals. This method offers personalised therapy and applies ML to analyse data. Robinson et al. [53] explore using AI embedded in smartwatches for emotional analysis and monitoring.

User-centric Emotion Monitor and Analysis

User-provided data offers deeper insights into emotional states. Identifying the best data to use is challenging. Randomised Controlled Trials are frequently used for this purpose to minimise bias [45]. Social interactions and behavioural actions can be valuable indicators of individuals' emotions. Mele et al. [89] employ five measures to assess well-being: quality of life, personal and subjective satisfaction, physiological and social functionality, physiological well-being, self-esteem, and quantity and quality of social relations. Several works use smartphones to capture user input [44]. Shi et al. [78] explore smartphones as tools for emotion regulation through experience sampling and interviews.

3.6.3 Therapeutic Approaches for Emotional Well-being

Therapeutic approaches fall into three types: virtual therapy, behavioural therapy, and AI-enabled therapy.

Virtual Therapy

Virtual therapy presents an optimistic route to fostering emotional wellness, offering immersive and readily available environments for therapeutic interventions. Through the immersive nature of virtual reality, individuals can participate in tailored treatment methods such as exposure therapy and mindfulness practices in a supportive setting [37]. This approach not only enhances accessibility to mental health support but also promotes emotional strength and general well-being in a compassionate and empathetic manner. Lingam [46] discusses

an application linked with the smartwatch that endeavours to comprehend users' behavioural patterns and suggest techniques to improve their well-being. Interactions and conversations with the app assist users in gaining insights into their emotional well-being.

Behavioural Interventions

Behavioural interventions play a vital role in promoting emotional well-being by assisting individuals in modifying their behaviours and developing positive lifestyle habits. Cognitive Behavioural Therapy (CBT) provides tangible methods for managing stress and anxiety, empowering individuals to support their resilience and improve their overall sense of well-being in a compassionate and supportive manner.

One method to enhancing mental well-being involves biofeedback therapy, where individuals are taught to gain voluntary control over a natural bodily function by first monitoring its normal automatic responses [37].

Yunusova et al. [62] explore the application of evidence-based therapies such as CBT and Acceptance and Commitment Therapy for treating depression. They delve into ecological momentary interventions, which involve delivering timely interventions to reduce stress and negative emotions as they arise in real-time situations. Thapa et al. [84] focus on exploring how technology can further support mental well-being, including mobile apps, wearable devices, and social media. They discuss how IoT could aid in predicting mental health outcomes and guiding treatment choices, offering promising avenues for improving overall well-being.

AI-Enabled Healthcare Technologies

AI-driven therapeutic methods are transforming emotional well-being systems by offering customised interventions through empathetic technologies such as chatbots and ML algorithms. These developments provide accessible and scalable assistance, fostering emotional resilience and well-being in a compassionate and empathetic manner. The AI-driven digital therapeutics field employs methods that can shape patient-centric approaches in developing virtual chatbots and referral systems [38].

3.7 IoT-based Solutions for Depression

Section 3.6 presented three categories of solutions: IoT-based, IoT integrated with AI methods and user-centric approaches. Depression related IoT solutions feature the same advantages

and disadvantages we identified for IoT solutions related to emotional well-being in Section 3.6. Figure 3.6 shows the categorisation of depression-related works from our review.

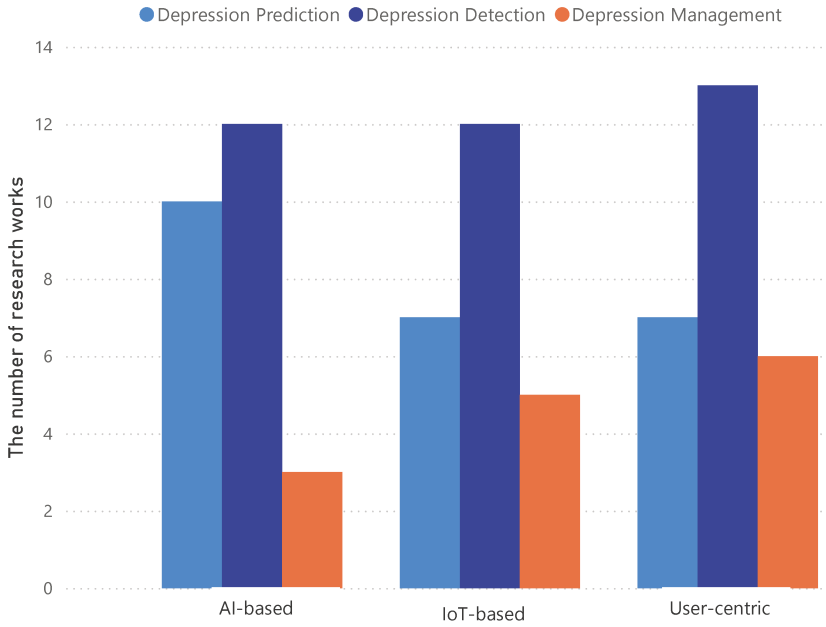


Fig. 3.7 Depression solutions based on type and method

Table 3.9 shows the works in each category and whether they perform prediction, detection, and/or management.

3.7.1 Predicting Depression

IoT-based Solutions

Moreira et al. [82] integrate IoT technology for real-time data collection in high-risk pregnancies to detect and predict post-partum depression. Moura et al. [68] focus on monitoring and predicting depressive states. They detect context-aware sociability patterns and changes in social behaviour through IoT sensors. IoT technology has been applied to mood monitoring, offering significant utility in evaluating and addressing psychological disorders [47].

AI-based Solutions

Srividya et al. [77] show that ML algorithms for clustering and classification could be expanded to predict different mental illnesses. They suggest incorporating physiological parameters to enhance the accuracy of predicting an individual's mental state. Kaur et al. [90] develop a fuzzy inference system for predicting levels of depression risk.

Table 3.9 Capabilities of depression-related IoT solutions

Method	Research Work	Predict	Detect	Manage
IoT-based Methods	Ou et al. [67]		✓	✓
	Moreira et al. [82]	✓	✓	
	Moura et al. [68]	✓	✓	
	Chen et al. [65]		✓	
	Thapa et al. [84]	✓	✓	
	Gupta et al. [56]		✓	
	Robinson et al. [53]	✓	✓	
	Jose et al. [70]	✓	✓	
	Gyrard et al. [88]	✓	✓	
	Mahdi et al. [69]		✓	
	Yang et al. [60]		✓	
AI-based Methods	Kumar et al. [55]	✓	✓	✓
	Srividya et al. [77]	✓		
	Kaur et al. [90]	✓	✓	
	Rajawat et al. [76]	✓	✓	
	Sinha et al. [83]		✓	
	Ali et al. [59]		✓	
	Machado et al. [57]	✓	✓	
	Kumar et al. [85]	✓	✓	
	Saffar et al. [39]		✓	
User-centric Methods	Mele et al. [89]	✓	✓	
	Lingam [46]			✓
	Yunusova et al. [62]	✓		✓
	Hardy et al. [71]		✓	

Physiological data is invaluable in predicting depression. Heart rate variability and body temperature, that have been used to predict major depressive disorders in [42]. Machado et al. [57] study the impact of stress, anxiety, and depression on the body's vital signs. Kumar et al. [55] propose an algorithm to forecast emotions utilising conversational text extracted from messaging applications and social media.

User-centric Approaches

Yunusova et al. [62] employ a personalised, predictive model of mental health utilising event mining to predict alterations in mental health, considering factors like disruptions in sleep patterns, challenges in social interactions, decreased physical activity, and reduced experience of positive emotions to anticipate heightened depression severity. Similarly, Mele et al. [89] employ ML to assess depression likelihood from patients' behaviour, enabling early interventions for targeted depression management.

3.7.2 Detecting Depression

IoT can help provide critical data such as facial expressions, voice analysis, and physiological information to recognise signs of depression [42].

IoT-based Solutions

Data from IoT devices, including Wireless Sensor Networks (WSN) and Wireless Body Area Networks (WBAN), is used to identify and treat depression through cognitive and motor rehabilitation methods in [42]. Gupta et al. [56] capture pulse and skin resistance via wearables to detect anxiety and depression. The IIFDD framework [65] uses IoMT-based remote monitoring and diagnostics as well as a feature extract module for detecting depression.

Thapa et al. [84] use speech landmark bigrams to detect depression from everyday smartphone speech. Moreover, Ou et al. [67] analyse behavioural data from IoT devices and utilise ML techniques to detect changes in depression levels. Robinson et al. [53] use wearable data-based metrics like physical activity, sleep duration, and self-reported social interactions to assess the severity of depressive symptoms. Jose et al. [70] use biosensors and phone sensors and integrate support apps and social media to detect and predict depression. Yang et al. [60] investigate the capabilities of a wearable social sensing platform to monitor participants' psychological conditions.

AI-based Solutions

Kaur et al. [90] combine Mamdani's fuzzy logic controller with a feed-forward multilayer neural network to detect depression. Moreover, employing multimodal detection systems that integrate computer vision, physiological signals, and IoT technology could offer valuable insights for detecting depression. Their multimodal ER model MEmoR [85] combines computer vision and IoT devices to capture visual and physiological data to detect depression.

Several studies use sentiment analysis on text, such as social media data, for detecting depression [37, 59]. Sinha et al. [83] develop large-scale emotion analysis and surveillance technology for identifying depression symptoms in student cohorts. Rajawat et al. [76] use AI-enabled robots for prolonged observation of user behaviour to detect depression. Saffar et al. [39] note that textual emotion detection is widely used in health and medicine, particularly for identifying depression.

User-centric Approaches

Questionnaires, such as the Beck Depression Inventory (BDI), are vital tools for gathering user-based data [60]. Kumar et al. [55] assess a person's baseline depressive state by gathering data from various sources. Hardy et al. [71] present an emotion questionnaire containing robust questions from reputable organisations.

3.7.3 Managing Depression with IoT

Lingam [46] explore using smartwatches and associated applications for depression monitoring. Their app facilitates data sharing with healthcare professionals to gain insights into individuals' conditions and provides suggestions for improving their well-being based on behaviour patterns. Kumar et al. [55] compare current depression levels with target levels to determine the appropriate intervention, including providing suggestions and/or using smart lighting to modulate mood. Ou et al. [67] deploy wireless sensors to capture activity patterns, which are fed into ML models to predict and detect changes in depression levels. Robinson et al. [53] underscore the role of sensors in monitoring mental status and providing timely and accessible data to various professionals and users.

3.8 Discussion

3.8.1 Data in Technological Emotional Well-being and Depression-Related Solutions

Data forms the backbone of well-being and healthcare systems, crucial to their effectiveness. When it comes to emotional well-being, capturing and analysing diverse types of data can improve effectiveness.

Building on the foundations covered in previous sections (see Sections 3.4 and 3.5), this study highlights how the integration of diverse types of data, data collection and analysis methods in emotional well-being and depression systems differs from prior reviews. Unlike previous studies that primarily focused on technologies, data collection, or analysis methods, our study provides a holistic understanding by combining the key data-related aspects of IoT-based solutions.

The most useful types of data for enhancing emotional well-being and managing depression through IoT are:

1. **Self-reported data:** Subjective information from individuals about mental and physical well-being.

2. **Physiological data:** Data related to physiological aspects of a user's body, such as EEG and heart rate.

3. **Behavioural data:** Data related to users' behaviour, such as interactions on social media.

4. **Environmental data:** Data from other IoT sensors representing environmental conditions, such as ambient temperature, humidity, etc.

Figure 3.8 illustrates the frequency of each type of data used in the solutions we reviewed. Table 3.10 shows that physiological data remains the most widely used in IoT-based solutions for emotional well-being and depression, consistent with trends identified in previous reviews. However, our study reveals a growing interest in integrating environmental data, which is a relatively new area. This shift underscores a broader understanding of how external factors contribute to mental health, suggesting new avenues for enhancing interventions and potential for addressing depression more comprehensively.

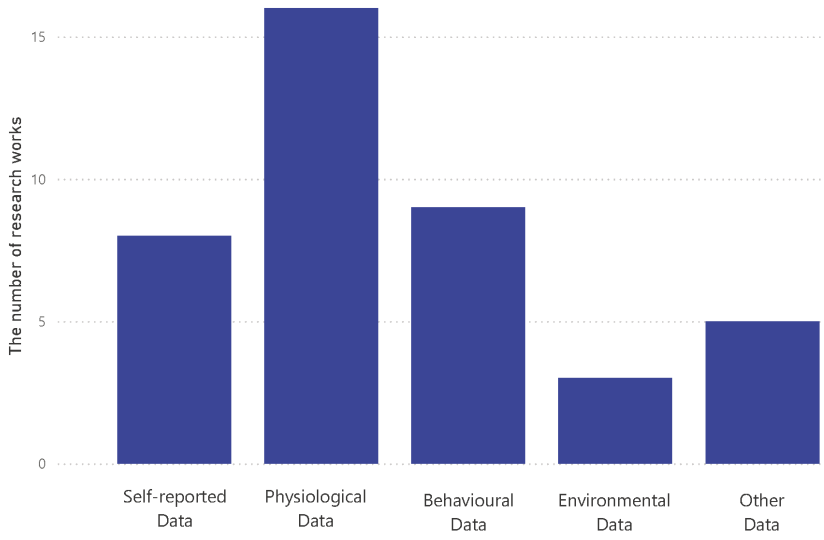


Fig. 3.8 Data in IoT-based emotional well-being and depression systems

Environmental data can play a crucial role in understanding and managing emotional well-being and depression. Critical data such as temperature, humidity, and air quality, provide objective information on the physical conditions that may influence mental health. Cameras and other environmental sensors can provide valuable insights into individuals' surroundings and behaviours. Visual cues such as facial expressions, body language, and interactions with the environment offer valuable contextual information for assessing and supporting emotional well-being.

Table 3.10 Types of data used in technological emotional well-being and depression systems

Type	Data	Reference
Physiological	Body temperature	[69], [70], [53]
	Respiration rate	[70]
	Heart rate	[70], [53], [55], [58]
	Blood pressure	[70]
	Pulse rate	[70], [69], [56]
	EDA	[53]
	GSR	[53], [56]
	Physical activity	[53]
	ECG	[84]
	EEG	[53], [91]
	Other signs	[60], [82], [85], [57], [63], [71]
Environmental	Ambient data	[61]
	IoRT data	[88]
	Smart home data	[66]
Behavioural	Text, audio, image and video	[76], [65], [84], [56], [90], [83]
	Activities log	[62]
	Online activity records	[68], [55]
	Other behavioural data	[60]
Self-reported	Survey data	[67], [77], [71], [78]
	Interviews data	[89]
	Other self-reported data	[62], [56]
Other	Datasets	[59], [82]

Figure 3.3 indicates that the growth of the IoT and sensors has seen these used more extensively in data collection. IoT technology can be seamlessly integrated into users' everyday environments, thus promoting better adherence. By embedding IoT devices and making the technology less visible, these solutions can reduce the burden on users and foster continuous engagement in their care. Interestingly, despite the advancements in IoT technology, our studies reveal that manual methods like questionnaires continue to play a critical role in capturing nuanced emotional states. These findings contrast with the diminishing emphasis on such methods in some reviews, highlighting their enduring relevance, especially when integrated with new technologies.

The widespread adoption of mobile devices in contemporary society warrants further exploration of their potential for research. In contrast to earlier studies that primarily explored the passive data collection capabilities of mobile devices, our findings suggest a dual role: not only can these devices collect data through embedded sensors, but they also actively engage users in providing contextual information. This dual approach enriches the data quality and offers a more comprehensive view of daily activities and behaviours.

Preprocessing data is vital for emotional well-being and depression support systems. This involves cleaning, organising, and transforming raw data gathered from different sources into usable information, ensuring accuracy and reliability in subsequent analyses. Our review identifies a significant gap in the literature regarding the preprocessing of data from IoT sensors, an area that is glaringly missing from most studies we reviewed. This underscores the necessity of developing standardised preprocessing protocols to ensure data accuracy and reliability, a critical step that could significantly impact future research outcomes in this field.

3.8.2 Effectiveness Parameters in Emotional Well-being and Depression Solutions

In healthcare systems, particularly within mental health and emotional well-being services, solutions must effectively address the needs of individuals facing various challenges. While standards like Health Level 7 (HL7) [92] outline some effectiveness parameters, our study proposes expanding these criteria to include emerging technologies and types of data not fully covered in current standards. This expansion could lead to more robust and adaptive healthcare systems tailored to the complex needs of mental health services. For example, regular tracking of patients' emotional states or subjective well-being, often through self-reports or digital tools, may require specialised data beyond regular standards, as specific emotional metrics and continuous mood tracking often need customisation.

Accurate and reliable data collection is critical in IoT-based emotional well-being and depression-related systems. Issues like connectivity problems, device malfunctions, or incorrect sensor placement can lead to incomplete or noisy data, impacting clinical decision-making. Robust data validation, quality assurance processes, and continuous sensor calibration and monitoring are essential to prevent inaccuracies that could result in misdiagnoses or ineffective treatments. User compliance and engagement are also significant challenges. Patients may only sometimes use IoT devices due to discomfort, technical difficulties, or privacy concerns. To address this, designing user-friendly device interfaces, providing education and training, and incorporating behavioural reinforcement mechanisms like reminders or incentives are all essential.

Another challenge is interoperability. IoT systems must integrate seamlessly with existing healthcare infrastructure, such as electronic health records, where data silos and varying standards often hinder information exchange. Moreover, privacy and security concerns around handling sensitive health data require secure transmission, encryption, and adherence to regulations. Finally, scalability is crucial for widespread adoption, particularly for smaller

clinics or underserved regions. Financial constraints can limit adoption, so solutions must be cost-effective and scalable.

3.8.3 Depression Management

A significant issue within the realm of mental healthcare lies in effectively addressing and supporting individuals diagnosed with depression. Previous reviews have largely concentrated on the detection and prediction of depression, as noted in Section 3.7. In contrast, our study shifts the focus towards the active management of depression, emphasising the development of personalised care plans that address the unique needs of individuals. This approach represents a significant evaluation from earlier works, which have often overlooked the ongoing support required for effective long-term management.

Furthermore, personalised care plans that consider factors like symptom severity, underlying causes, concurrent medical conditions, and individual preferences are vital for optimising outcomes in depression management. This approach ensures that individuals receive comprehensive support tailored to their unique circumstances, fostering better long-term outcomes and quality of life. IoT offers personalised care by leveraging the ability to adapt to individual users' unique contexts and needs.

By shifting focus towards actively managing depression, healthcare systems can better address the diverse needs of individuals living with this condition, ultimately improving overall mental health outcomes.

3.9 Conclusion

This article explored the evolving landscape of IoT technologies in monitoring, analysing, and enhancing emotional well-being and depression. Our analysis identified a multifaceted approach towards understanding and managing depression, leveraging a wide array of data collection methods—from sensor-based and user-centric to using datasets. Integrating IoT devices and AI algorithms has shown promise in providing personalised, real-time insights into individuals' emotional states, potentially revolutionising how emotional well-being is approached in healthcare and everyday settings.

One of the critical revelations of this review is the significant potential that IoT technologies hold in enhancing the granularity and immediacy of emotional well-being monitoring. Wearable sensors, mobile applications, and intelligent environments facilitate continuous data collection and ensure data richness and contextuality, which is crucial for accurate emotion detection and analysis. Furthermore, AI-based methods, including ML and NLP,

have demonstrated efficacy in deciphering complex emotional patterns, offering pathways to predictive analytics in mental health.

However, the journey toward fully harnessing the capabilities of IoT and AI in emotional well-being is fraught with challenges. Issues surrounding data privacy, security, and the ethical implications of continuous monitoring remain paramount. Additionally, the heterogeneity of data and sources necessitates advanced data integration and analysis frameworks to ensure coherent and actionable insights.

Looking forward, the convergence of IoT and AI technologies presents a transformative opportunity for mental health. Future research should aim at refining these technologies to address current limitations, exploring novel applications, and ultimately facilitating their seamless integration into healthcare systems and daily life. The ultimate goal is to create an empathetic, intelligent, and interconnected ecosystem that anticipates and responds to emotional distress and fosters overall emotional resilience and well-being.

In conclusion, the synergistic combination of IoT data solutions and AI promises to reshape the landscape of depression management and emotional well-being. By pushing the boundaries of current technological capabilities and addressing the associated ethical and practical challenges, we can pave the way for innovative, impactful, and empathetic solutions, offering hope and support to those navigating the complexities of mental health.

Future directions include exploring the impact of various categories of data for managing emotional well-being and depression. Future research should prioritise the development of predication and active management solutions specifically designed for individuals already diagnosed with depression, an area underexplored in prior studies. It is also essential to consider the ethical implications of integrating IoT and AI in mental health interventions. Data privacy, informed consent, and the potential for algorithmic bias must be addressed to protect individuals and ensure equity and responsible use. Another interesting direction is to study factors to develop effective IoT-based solutions for other mental health disorders, such as stress and anxiety. A generalised review of AI-based solutions is also a promising area for future research.

Chapter 4

Prelude to Contextualising Environmental Data in IoT-Based Mental Health Monitoring

4.1 Introduction

Chapter 5 focuses on an important contextual layer in IoT-based mental health monitoring: the physical environment. In this thesis, context refers to the surrounding conditions that help interpret behavioural, physiological, or self-reported signals, rather than treating them as isolated indicators. For Chapter 5, this context is specifically environmental: the temperature and humidity conditions of enclosed spaces that may be associated with comfort, stress, sleep, concentration, and emotional well-being [93–96]. This focus responds to a limitation identified in Chapters 1–3. The earlier chapters showed that many IoT-based mental health systems focus on isolated physiological, behavioural, or self-reported signals, with limited attention to the contextual factors that may shape or explain those signals. For example, a behavioural change such as reduced activity or disturbed sleep may be interpreted differently depending on the person's environment, routine, and living conditions. Environmental data can provide useful contextual information for interpreting mental health-related patterns.

The purpose of Chapter 5 is to examine environmental sensing as a contextual sensing exploration within a broader IoT-based monitoring framework. The chapter investigates how temperature and humidity data in enclosed environments can be collected, processed, classified, and used to support a more contextual interpretation of stress-related environmental conditions. This exploration shows how environmental data may function as one support-

ing contextual layer within a more comprehensive, personalised, and adaptive IoT-based approach.

4.2 Role of Chapter 5 in the Thesis

Chapter 5 represents the first technical study after the systematic review. While the review identified gaps in the literature, this chapter begins to explore how one of those gaps can be examined through a technical study. Specifically, it examines how environmental data can be integrated into IoT-based mental health monitoring as a contextual input. This is important because many IoT-based health and well-being systems focus on physiological, behavioural, or self-reported data, while environmental context is often treated as secondary or omitted entirely [97, 98].

Chapter 5 plays an important role in advancing the thesis. Rather than moving directly from the systematic review to the final architecture, the thesis first examines selected dimensions that are important for a more complete IoT-based depression management system. Chapter 5 contributes to this progression by investigating contextual sensing through environmental data. The following chapters build on this foundation by considering additional dimensions, including user engagement, system integration, safety, personalisation, and expert-informed refinement. In this way, Chapter 5 acts as a bridge between the literature-based understanding of the field and the architecture-oriented contribution developed later in the thesis. It shows that environmental context is not only a theoretical consideration, but can also be operationalised as a technical layer within an IoT-based monitoring framework. Therefore, the chapter helps justify why the later architecture includes contextual sensing as part of its broader design.

Chapter 5 contributes to the thesis in three main ways:

1. Positioning environmental context as a relevant supporting layer in IoT-based mental health monitoring. This supports the broader thesis argument that IoT-based depression management systems should support interpretation of data in relation to the person and their surrounding context, rather than treating sensor data as isolated indicators.

2. Providing a data-driven exploration of environmental sensing. This chapter uses temperature and humidity datasets from enclosed environments and applies machine learning techniques to classify environmental conditions using predefined comfort- and stress-related categories. This demonstrates how ambient data can be processed and analysed within a monitoring pipeline. Previous studies have reported associations between temperature, humidity, heat-related environmental conditions, and mood,

emotional well-being, or mental health outcomes [96, 99]. Building on this evidence, Chapter 5 explores how such factors can be operationalised as contextual inputs within an IoT-based data-processing framework.

3. Highlighting the need for personalisation and careful interpretation. The results show that environmental data can support classification of stress-related environmental conditions, but they also reveal limitations, including reliance on temperature and humidity as the main environmental variables, use of predefined thresholds, class imbalance, absence of depression-specific outcome data, and lack of direct integration with behavioural, physiological, and self-reported mental health data. These limitations reinforce a key argument of this thesis: technical prediction alone is not sufficient for mental health monitoring. A safe and useful system must also consider individual differences, uncertainty, context, and the risk of misinterpretation.

4.2.1 Interpretation and limitations of environmental stress labels

An important distinction concerns the interpretation of the stress categories used in Chapter 5. The three datasets analysed in the study contain environmental measurements of temperature and humidity. The five categories (Very Low, Low, Medium, High, and Very High) were constructed using predefined temperature and humidity thresholds informed by findings from previous studies and domain-specific knowledge, as presented in Table 5.1 and Sections 5.4.2 and 5.4.3. These categories represent literature-informed proxy labels assigned to environmental conditions [100, 96, 101–103, 99]. The term *stress level*, as used in Chapter 5, should consequently be interpreted as an environmental stress-related classification rather than a direct measure of an individual's experienced stress or depression severity.

The thresholding approach assumes that environmental conditions associated in the literature with thermal discomfort or stress-related responses can provide a useful basis for categorising environmental observations. However, this assumption does not establish a direct or universal relationship between specific temperature–humidity combinations and individual mental health states. Responses to environmental conditions may vary according to personal sensitivity, acclimatisation, duration of exposure, health status, and other contextual factors.

Consequently, the machine learning results reported in Chapter 5 should be interpreted as measures of the models' ability to classify environmental observations according to the constructed labels. Therefore, the contribution of Chapter 5 is a technical exploration of environmental data processing and classification that supports the inclusion of environmental context within IoTMindCare. Future empirical work would be required to examine how

such contextual variables relate to independently measured mental health outcomes and how they should be combined with behavioural, physiological, and self-reported information for personalised interpretation.

4.3 Connection to Subsequent Chapters

The following chapters build on the contextual foundation established in Chapter 5. While Chapter 5 focuses on environmental context, the next stage of the thesis examines user interaction and engagement through gamified AI-IoT-based monitoring and feedback. This progression is important because contextual sensing alone does not ensure sustained use, self-awareness, or meaningful user participation. Therefore, the subsequent chapter extends the thesis from passive environmental sensing toward more interactive and user-facing forms of IoT-based mental health support.

The findings from Chapter 5 also inform the development of the *IoTMindCare* architecture in Chapter 9. Specifically, the chapter indicates that an environmental sensing layer involves continuous data collection, preprocessing, and classification, while its integration into a broader monitoring system would also require personalisation and combination with other data streams. However, Chapter 5 also identifies several limitations that must be addressed at the architectural level. These include the reliance on temperature and humidity as the main environmental variables, the use of predefined stress-related thresholds, class imbalance in the classification results, and the absence of direct integration with behavioural, physiological, and self-reported mental health data. Taken together, these findings and limitations support treating environmental sensing as one component of a broader context-aware monitoring system rather than as an independent indicator of mental health status. Consequently, within the *IoTMindCare* architecture, environmental data are integrated alongside behavioural, physiological, self-reported, and interaction-based information to support a more comprehensive approach to depression management.

Chapter 5

Integrating Environmental Data for Mental Health Monitoring: A Data-Driven IoT-Based Approach

5.1 Abstract

Mental health disorders constitute a significant global challenge, compounded by the limitations of traditional management approaches that rely heavily on subjective self-reports and infrequent professional evaluations. This study presents a groundbreaking IoT-based system that integrates big data analytics, fuzzy logic, and machine learning to revolutionise mental health monitoring. In contrast to existing solutions, the proposed system uniquely incorporates environmental factors, such as temperature and humidity in enclosed spaces–critical yet often overlooked contributors to emotional well-being. By leveraging IoT devices to collect and process large-scale ambient data, the system provides real-time classification and personalised visualisation tailored to individual sensitivity profiles. Preliminary results reveal high accuracy, scalability, and the potential to generate actionable insights, creating dynamic feedback loops for continuous improvement. This innovative approach bridges the gap between environmental conditions and mental healthcare, promoting a transformative shift from reactive to proactive care and laying the groundwork for predictive environmental health systems.

Publication Note

This chapter is based on the published article:

Zamani, S., Nguyen, M., & Sinha, R. (2025). Integrating environmental data for mental health monitoring: A data-driven IoT-based approach. Applied Sciences, 15(2), 912.

The published article is available online via https://doi.org/10.3390/app15020912.

The remainder of this chapter is based on the published article, with minor formatting adjustments made to align it with the structure and argument of this thesis.

5.2 Introduction

Mental health disorders constitute a critical global health challenge, impacting millions of lives and imposing significant social and economic burdens. The World Health Organization (WHO) identifies mental health conditions such as anxiety, depression, and stress-related disorders as leading causes of disability worldwide [104]. Despite increasing awareness, existing mental health management strategies remain hindered by reliance on subjective self-reports and periodic professional evaluations [105], which often lack objectivity, immediacy, and accessibility. These shortcomings severely limit the delivery of effective, personalised interventions, particularly in resource-limited environments.

5.2.1 Mental Health and Environmental Factors

Recent advancements in emotional well-being monitoring systems have begun to address these challenges by utilising data-driven techniques to identify emotional states [106]. Many of these systems leverage wearable sensors, smartphones, and other digital platforms to gather physiological and behavioural data, offering insights into individuals' mental health [107]. However, these solutions primarily focus on internal factors, such as heart rate variability and sleep patterns, while neglecting the impact of external environmental conditions, including temperature, humidity, air quality, and lighting. Research has shown that such environmental factors significantly influence emotional well-being and mental health, particularly in enclosed spaces where people spend a substantial portion of their time [93, 94]. Ignoring these factors limits the comprehensiveness and efficacy of current monitoring solutions. Studies have shown that temperature, humidity, and other ambient factors can profoundly affect mental health [108]. However, they remain underutilised in treatment and prevention strategies.

5.2.2 The Role of Big Data and IoT in Mental Health Management

In this context, the Internet of Things (IoT) has emerged as a transformative technology capable of addressing mental health monitoring gaps. IoT devices equipped with environmental

sensors can continuously collect vast amounts of real-time data, creating new opportunities for monitoring the interplay between environmental conditions and mental health. The application of IoT in health monitoring is increasingly prevalent, with several studies exploring its role in managing chronic conditions, sleep disorders, and stress levels [109, 97, 98]. When combined with big data analytics, fuzzy logic, and machine learning (ML) techniques, IoT systems can process and analyse these datasets to uncover meaningful patterns, provide actionable insights, and deliver personalised recommendations. With the proliferation of IoT devices, environmental data can now be collected continuously at an unprecedented scale. Big data analytics offers a powerful framework to process, analyse, and extract actionable insights from these data.

This paper proposes a novel IoT-based system for mental health monitoring that integrates these technologies to address critical shortcomings in traditional and emerging methods. The system aims to offer real-time classification, visualisation, and predictive analytics tailored to individual needs by capturing and analysing ambient environmental data alongside personal sensitivity profiles. The proposed approach enhances the accuracy and relevance of mental health monitoring. It fosters a shift from reactive to proactive care by enabling early detection of adverse conditions and their impact on emotional well-being.

This study aims to address the following questions:

1. How can big data and IoT technologies improve environmental monitoring for mental health management?

2. What predictive capabilities can be introduced to anticipate adverse conditions?

3. How can these predictions be personalised to account for individual sensitivities?

The remainder of this paper is organised as follows: Section 5.3 reviews related work in IoT-based mental health monitoring and the role of environmental factors. Section 5.4 outlines the design and architecture of the proposed system, detailing the integration of IoT devices, big data analytics, fuzzy logic, and machine learning. Section 5.5 presents initial experimental results, highlighting the system's performance in terms of accuracy, scalability, and real-world applicability. Section 5.6 discusses the implications of the findings, including their potential to empower mental healthcare providers and individuals. Finally, Section 5.7 concludes the paper, summarising contributions and outlining directions for future research.

5.3 Related Work

5.3.1 Environmental Factors and Mental Health

Environmental factors play a significant role in influencing mental health outcomes. Various studies have highlighted the impact of physical surroundings, such as urban environments, pollution, noise, and climate change, on psychological well-being. Urbanisation, for instance, has been linked to increased stress levels and higher risks of anxiety and depression due to factors like overcrowding, lack of green spaces, and heightened exposure to noise pollution [110, 111]. Air pollution, particularly exposure to fine particulate matter (PM2.5), has been associated with cognitive decline and the exacerbation of psychiatric disorders [112]. Similarly, natural disasters and extreme weather events, which are intensifying due to climate change, have been shown to contribute to post-traumatic stress disorder (PTSD), depression, and anxiety disorders [95].

Understanding these relationships underscores the need for interdisciplinary approaches combining public health, environmental science, and urban planning to address the mental health challenges associated with environmental factors.

5.3.2 Recent Advancements in Predictive Modelling

Recent advancements in predictive modelling have significantly contributed to environmental monitoring and mental health fields. Among these, two of the most notable developments are the application of cognitive computing to predict the flow status of flexible rectifiers and the predictive modeling of flexible electrohydrodynamic (EHD) pumps using Kolmogorov–Arnold Networks (KAN).

The first advancement leverages cognitive models to enhance the accuracy of stress prediction systems by integrating environmental and physiological data. Peng et al. [113] explore this methodology, demonstrating its potential in adaptive stress detection systems.

The second advancement involves the use of KAN, which employs learnable spline-based activation functions to model complex nonlinear relationships. This approach offers improved predictive accuracy over traditional models. Peng et al. [114] highlight KAN's effectiveness in capturing intricate interactions between environmental factors and individual responses, which can be applied to refine stress detection models for better precision in predicting mental health outcomes.

These advancements underscore the growing role of intelligent, flexible systems in enhancing prediction accuracy within uncertain and dynamic environments, with promising applications in mental health prediction and stress detection.

5.3.3 Big Data in Healthcare

The integration of big data analytics in healthcare has transformed the way medical professionals diagnose, treat, and manage diseases. With the proliferation of electronic health records (EHRs), wearable devices, and other health-related technologies, an unprecedented volume of data are now available for analysis [115]. Big data enable the identification of patterns and trends in patient populations, facilitating predictive analytics and personalised medicine. For example, machine learning algorithms can analyse large datasets to predict disease outbreaks, optimise treatment plans, and improve patient outcomes [116].

However, challenges such as data privacy, interoperability, and the need for advanced computational infrastructure must be addressed to fully harness the potential of big data in healthcare. As the field evolves, ethical considerations surrounding the use of patient data and the potential for algorithmic bias remain critical areas of focus [117].

5.3.4 Fuzzy Logic in Environmental Data Analysis

Fuzzy logic, a mathematical approach based on degrees of truth rather than binary logic, has emerged as a powerful tool for analysing complex and uncertain environmental data. Environmental systems often involve intricate interdependencies and imprecise information, making traditional analytical methods less effective [118]. Fuzzy logic provides a framework for dealing with such uncertainties, enabling researchers to model complex environmental phenomena more accurately. Applications of fuzzy logic in environmental data analysis include air quality monitoring, climate change modelling, and disaster risk assessment. For instance, fuzzy logic systems can evaluate air quality indices by integrating multiple parameters, such as pollutant concentrations, humidity, and temperature, to produce comprehensive assessments [119].

The adaptability of fuzzy logic allows for the incorporation of expert knowledge and subjective judgments, making it particularly useful for decision-making in environmental management. Despite its advantages, the implementation of fuzzy logic requires careful design and validation to ensure reliability and accuracy. As environmental challenges become increasingly complex, the role of fuzzy logic in providing actionable insights is expected to grow significantly.

5.4 Methodology

This section outlines the methodology employed in the proposed IoT-based mental health monitoring system, which leverages big data, machine learning, and environmental sensors to

monitor temperature and humidity levels in enclosed settings for managing the mental health status of residents. The system's objective is to identify correlations between environmental parameters (specifically temperature and humidity) and emotional states (particularly stress level), as well as to develop predictive models for early intervention.

5.4.1 Datasets

For this study, three publicly available datasets containing temperature and humidity measurements from enclosed environments were selected. These datasets were chosen because they offer a variety of environmental conditions and span multiple settings, providing a diverse range of data points for analysis.

1. **Correlating Indoor and Outdoor Temperature and Humidity in a Sample of Buildings in Tropical Climates** [120]—This dataset includes absolute humidity, relative humidity, and temperature measurements for indoor and outdoor environments across 35 buildings in eight tropical cities. We extracted Tmp and RH data from each building to create a consolidated dataset, referred to as *dataset1*, for our subsequent analyses.

2. **Indoor Temperature and Relative Humidity Dataset** [121]—This dataset comprises indoor temperature and relative humidity measurements collected over a month using twelve Xiaomi Mijia sensors in Medellín, Colombia. We consolidated all sensor data into *dataset2* for further analysis.

3. **Temperature and Humidity Dataset** [122]—This dataset includes temperature and humidity measurements for three regions, including Shenyang, Beijing, and Guangzhou. The data from all three cities have been combined into *dataset3* for subsequent analysis.

Each dataset contains time-stamped records of temperature and humidity levels, which were merged into a unified database for analysis. The primary objective was to analyse the relationship between these environmental variables and the emotional well-being of individuals within these environments.

5.4.2 Data Collection and Preprocessing

Data were collected using IoT-enabled devices, including temperature and humidity sensors, which transmitted real-time readings to central servers. These IoT devices were distributed across the environments. They collected data at different intervals depending on the dataset, as described in Section 5.4.1.

Environmental factors, particularly temperature and humidity, significantly influence emotional well-being and stress levels. Extreme temperatures and high humidity have been linked to increased stress, discomfort, and negative emotional states, while moderate conditions are generally more conducive to emotional balance and reduced stress [100]. Research indicates that high temperatures can negatively affect mental health, leading to increased hospitalisation for mental illness [96]. Moreover, high humidity levels, particularly when combined with elevated temperatures, can exacerbate stress and discomfort [101]. Conversely, moderate temperatures and humidity levels are associated with better emotional well-being and reduced stress [102]. On the other hand, exposure to cold temperatures can lead to increased stress and negative mental health outcomes. For instance, individuals living in cold homes are at a higher risk of severe mental distress [103].

High humidity levels, particularly when combined with elevated temperatures, can exacerbate stress and discomfort. Florido et al. [99] found that relative humidity showed a significant correlation with suicide rates, indicating that both temperature and humidity are important factors in mental health.

Based on findings from prior studies, the relationship between temperature, humidity, and stress levels can be systematically analysed and categorised into distinct categories, as outlined in Table 5.1. These categories provide a structured framework for associating environmental conditions with stress levels, serving as a foundation for real-time stress monitoring and personalised mental health interventions. The thresholds used for categorising stress levels were informed by a combination of empirical studies and domain-specific knowledge.

Table 5.1 Categorisation of Stress Levels Based on Temperature and Humidity.

Stress Level	**Temperature Range**	**Humidity Range**
Very High	Below -20 °C or above 30 °C	Above 80%
High	Below -10 °C or 31 °C to 35 °C	70% to 80%
Medium	-10 °C to -1 °C or 26 °C to 30 °C	50% to 70%
Low	0 °C to 15 °C	30% to 50%
Very Low	16 °C to 25 °C	40% to 60%

These thresholds represent general boundaries derived from aggregated data, highlighting typical stress responses to specific temperature and humidity levels. However, it is crucial to recognise that these categories may require personalisation to accurately reflect individual experiences and environmental contexts. Factors such as acclimatisation to local climates, cultural variations in comfort levels, and individual health conditions significantly influence how environmental factors impact stress and emotional well-being. Future implementations

of these categorisations should consider adaptive mechanisms to fine-tune thresholds based on personal and contextual variables, ensuring more accurate and meaningful assessments.

The preprocessing steps involved the following:

- **Data Cleaning**: Missing data points were handled by using interpolation methods based on surrounding values. Outliers were detected using z-scores and were removed if they deviated significantly from the rest of the dataset.
- **Normalisation**: Temperature and humidity readings were normalised to a standard scale to ensure uniformity and improve the performance of machine learning models.
- **Data Labelling**: Emotional states were labelled based on the literature, as mentioned above. This labelling helped establish the ground truth for training the machine learning models.

5.4.3 Machine Learning and Big Data Integration

Given the volume and complexity of the data from three different datasets mentioned in Section 5.4.1, we used big data tools for distributed data processing and machine learning. The following steps were undertaken:

- **Data Processing:** The data were processed using standard machine learning and deep learning techniques in Python. The datasets were loaded using pandas and then cleaned and preprocessed through normalisation and feature scaling. We applied z-score normalisation to temperature and humidity values and handled missing data through interpolation. The final dataset was then labelled based on predefined stress levels derived from temperature and humidity conditions. Models were trained and evaluated on a 70–30% split for training and testing data.
- **Feature Selection:** A combination of statistical analysis and domain knowledge was used to identify the most relevant features for predicting emotional states. The selected features included temperature (T), humidity (RH), and historical temperature and humidity values, which were used to classify stress levels into five categories: very low, low, medium, high, and very high.
- **Model Training:** Several machine learning algorithms were implemented, including:
    1. **Logistic Regression (LR):** LR was chosen as a baseline model due to its simplicity, interpretability, and efficiency for binary and multi-class classification tasks. It is particularly effective for datasets where the relationship between features

(temperature and humidity) and the target (stress levels) is approximately linear, making it a natural choice for initial model comparison.

2. **Support Vector Machines (SVM):** SVM was selected for its ability to perform well in high-dimensional spaces and its robustness in handling both linear and non-linear decision boundaries. The Radial Basis Function (RBF) kernel was used to capture the complex relationships between temperature, humidity, and stress levels, especially when the data are not linearly separable. This makes SVM a powerful tool for classification problems where accuracy is critical.
3. **Long Short-Term Memory (LSTM) Networks:** LSTM was chosen to model the sequential dependencies between historical temperature and humidity values. As a type of Recurrent Neural Network (RNN), LSTM is designed to capture long-term dependencies, which is critical for predicting stress levels based on time-series data. This ability to account for temporal patterns makes LSTM ideal for scenarios where emotional states may evolve over time, driven by the continuous fluctuations in temperature and humidity.

Each model was trained using a 70–30% split for training and testing data. We performed hyperparameter tuning for each algorithm using grid search and cross-validation to optimise model performance.

- **Model Evaluation:** The models were evaluated using various performance metrics, including accuracy, precision, recall, and F1 score. We also assessed the ability of each model to predict stress levels based on historical temperature and humidity values, simulating a 30 min prediction window to enable early intervention in real-world scenarios.

5.4.4 System Architecture

This system is designed to enable continuous environmental data collection and real-time analysis for stress level prediction. The system architecture is illustrated in Figure 5.1, which provides a detailed representation of the integration between data collection, processing, and predictive analysis.

The system comprises the following key components:

- **IoT Devices:** These devices are equipped with sensors to continuously monitor temperature and humidity levels, which are critical environmental factors influencing emotional states. The real-time data collected by these sensors form the foundation for predicting stress levels by capturing fluctuations in environmental conditions.

- **Data Collection and Processing Platform:** This platform handles the acquisition and preprocessing of sensor data using Python libraries such as pandas and NumPy. The preprocessing workflow includes addressing missing values, normalising temperature and humidity, and calculating z-scores to standardise the data. This ensures the data are clean and ready for machine learning model training and real-time predictions.
- **Data Processing and Machine Learning Framework:** The processed data are utilised to train and evaluate multiple machine learning models, including Logistic Regression (LR), Support Vector Machines (SVM), and Long Short-Term Memory (LSTM) networks. Training is performed using a 70–30% train–test data split, with hyperparameter tuning applied to optimise model performance. These models are rigorously evaluated using metrics such as accuracy, precision, recall, and F1 score to ensure dependable predictions of stress levels.
- **Model Deployment:** The trained machine learning models are deployed to provide real-time predictions of emotional states based on live environmental data. The models categorise stress levels into five predefined categories, enabling dynamic assessment of a user's emotional well-being in response to current temperature and humidity conditions.

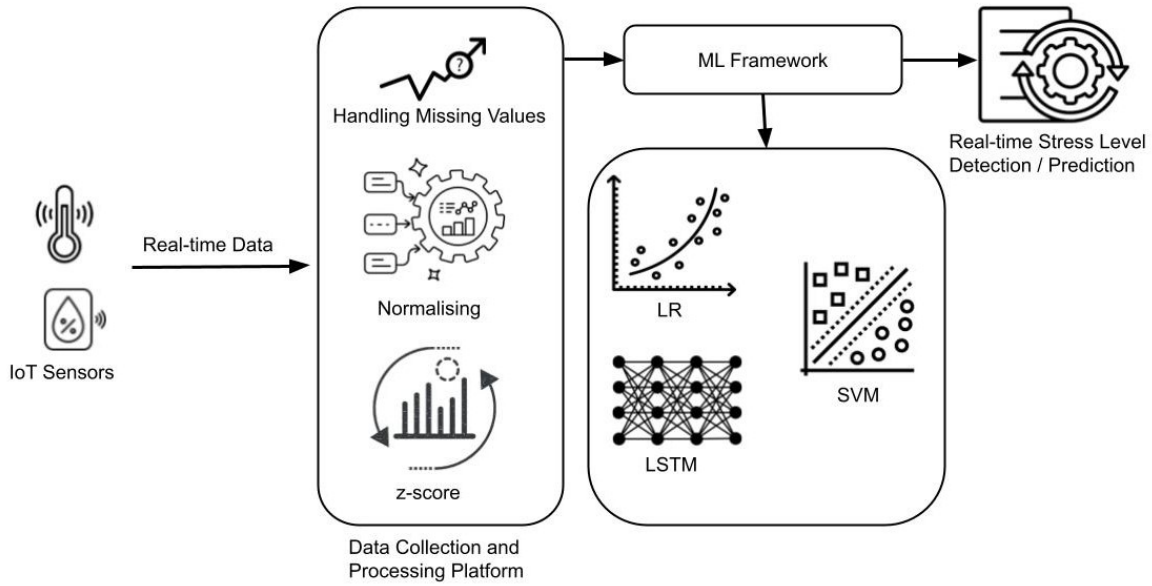


Fig. 5.1 System Architecture: A comprehensive framework integrating IoT devices, data processing platforms, and machine learning models for real-time stress monitoring and prediction.

5.5 Results and Experiment

This section presents the experimental results based on the three datasets, focusing on the analysis of temperature and humidity levels in enclosed environments and their relationship with emotional states. The following experiments evaluate the effectiveness of the proposed system using three machine learning models: Logistic Regression (LR), Support Vector Machine (SVM), and Long Short-Term Memory (LSTM).

5.5.1 Performance of Machine Learning Models

Table 5.2 provides a comprehensive summary of the performance of three machine learning models—Logistic Regression (LR), Support Vector Machines (SVM), and Long Short-Term Memory (LSTM) networks—in classifying stress levels based on environmental data. The evaluation metrics include accuracy, precision, recall, and F1 score, offering a detailed assessment of each model's effectiveness.

Table 5.2 Performance Metrics of Machine Learning Models for Stress Level Classification.

Model	**Accuracy (%)**	**Precision (%)**	**Recall (%)**	**F1 Score (%)**
LR	96.04	94.34	96.04	95.18
SVM	97.95	96.22	97.95	97.07
LSTM	70.03	49.00	70.03	56.65

The results highlight that the SVM model outperformed the others, achieving the highest accuracy (97.95%) and demonstrating balanced performance across all metrics. The Logistic Regression (LR) model also performed well, with a slightly lower accuracy of 96.95%. In contrast, the LSTM model showed moderate performance, with an accuracy of 70.03% but significantly lower precision (49%), likely due to the limited size of the training dataset and a lack of diverse temporal patterns in the environmental data. These findings suggest that while SVM is the most suitable for stress level classification in this context, the LSTM model's performance may be improved with further optimisation, such as hyperparameter tuning (e.g., adjusting learning rate, number of layers, and units per layer) and by using a more suitable dataset with stronger temporal dependencies. In addition, the lack of diverse temporal features in the dataset may have limited LSTM's ability to capture important patterns, potentially limiting its ability to improve prediction accuracy. This insight opens avenues for future work, where temporal models like LSTM might perform better with more data or refined feature engineering.

5.5.2 Predictive Accuracy

Figure 5.2 illustrates a comparative analysis of the predictive accuracy of the three machine learning models—Logistic Regression (LR), Support Vector Machines (SVM), and Long Short-Term Memory (LSTM) networks—based on their classification results. The chart highlights that the SVM model consistently outperformed LR and LSTM across all evaluation metrics, including precision, recall, and F1 score. This visual comparison validates the robustness and reliability of the SVM model for classifying stress levels using static environmental data.

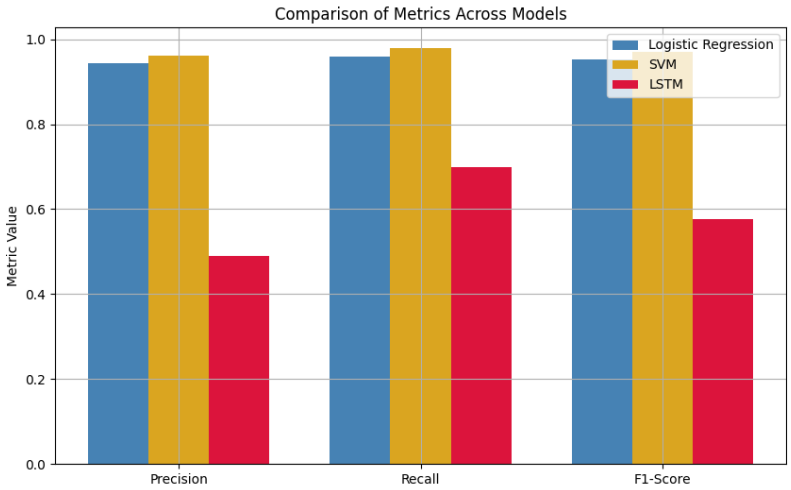


Fig. 5.2 Comparison of Precision, Recall, and F1 Score Across Machine Learning Models: SVM Outperforms LR and LSTM in Handling Static Environmental Data.

Each model's performance was further analysed using confusion matrices for the five stress level categories: Very Low, Low, Medium, High, and Very High. These matrices provided detailed insights into the models' classification capabilities and highlighted areas for potential improvement. The following subsections discuss the key findings from this evaluation in detail, focusing on each model's strengths and limitations in addressing the classification task.

5.5.3 Class Imbalance and Model Performance

Figures 5.3–5.5 illustrate the impact of class imbalance on the performance of the Logistic Regression (LR), Support Vector Machine (SVM), and Long Short-Term Memory (LSTM) models. The dataset is dominated by the "Very Low" and "Very High" classes, which skew model performance towards these categories. While all models demonstrate high accuracy

for the “Very Low” class, their ability to correctly classify the “Low”, “Medium”, and “High” classes is significantly limited.

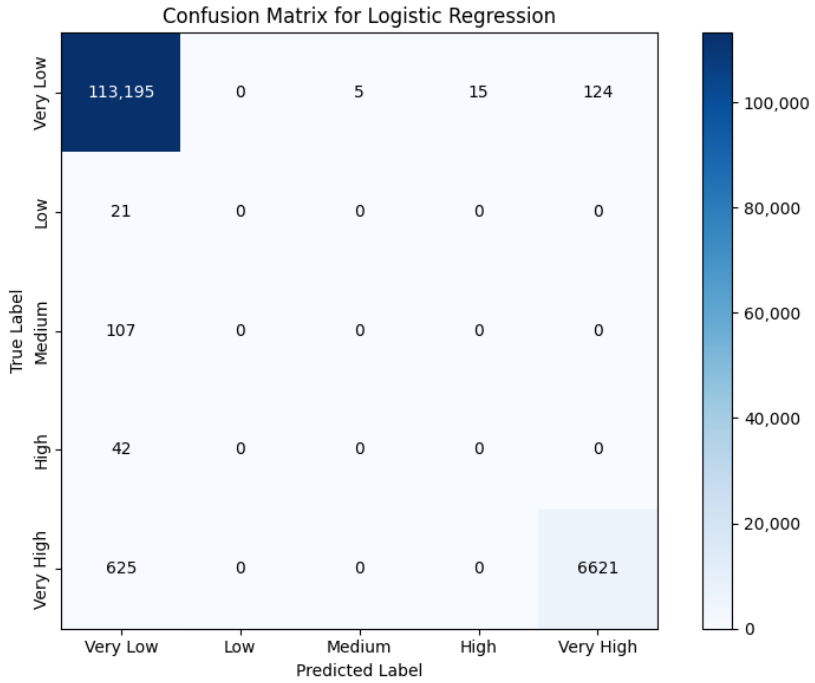


Fig. 5.3 Confusion Matrix for Logistic Regression (LR): The model performs well for dominant classes but struggles with mid-range categories.

Logistic Regression

As shown in Figure 5.3, the strengths and weaknesses of the LR model are as follows:

- **Strengths:** The LR model correctly predicts 113,195 instances for the “Very Low” class and 6621 instances for the “Very High” class, showcasing its strong performance on the dominant classes.
- **Weaknesses:** The model struggles with mid-range classes, misclassifying 625 instances of “Very High” as “Very Low” and 124 instances of “Very Low” as “Very High”. This indicates limited capacity to differentiate between these mid-range stress levels.

Support Vector Machine

As depicted in Figure 5.4, the SVM model exhibits the following strengths and weaknesses:

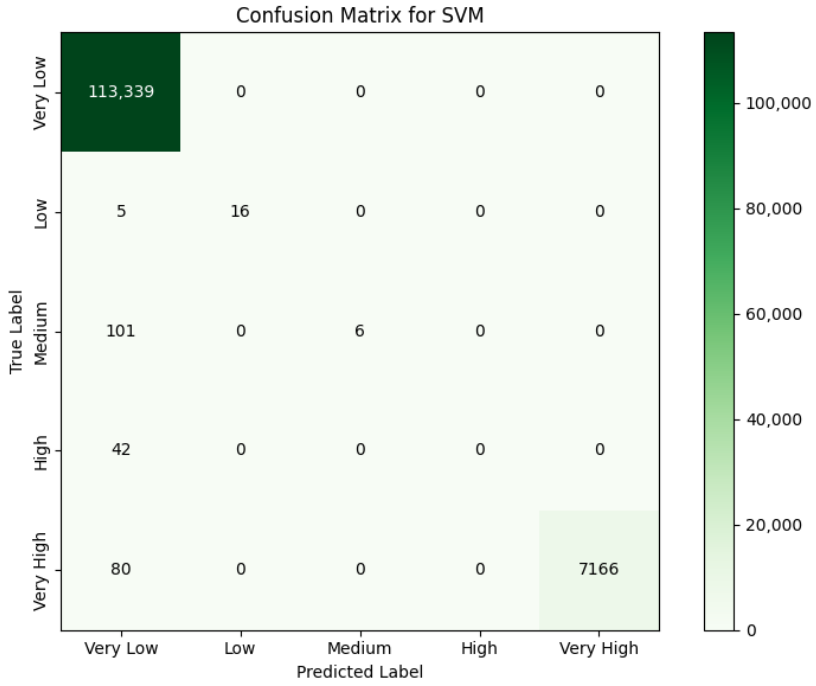


Fig. 5.4 Confusion Matrix for Support Vector Machine (SVM): The model shows improved accuracy for "Very High" and slight improvements for "Low" and "Medium" classes.

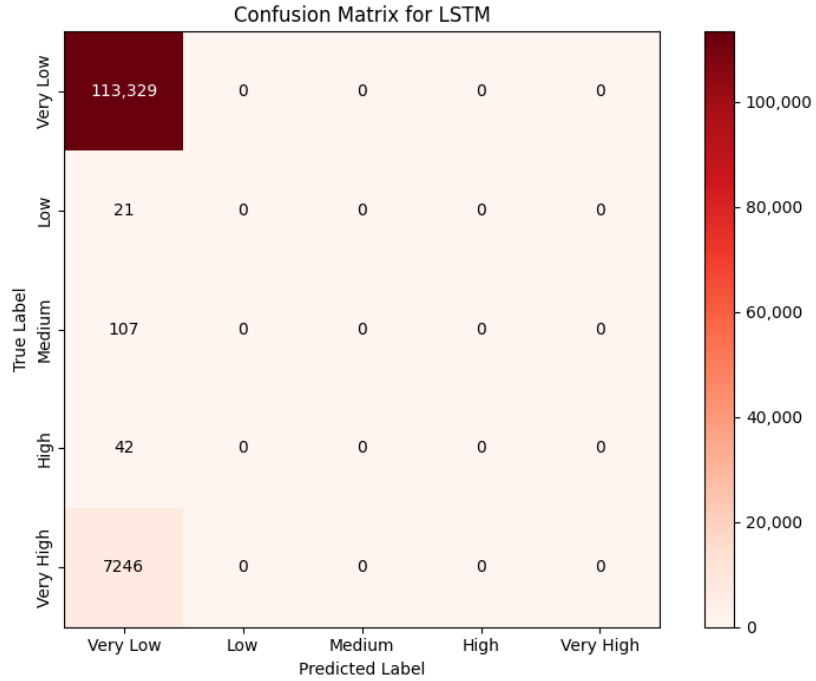


Fig. 5.5 Confusion Matrix for Long Short-Term Memory (LSTM): The model excels in predicting "Very High" but fails entirely for "Low", "Medium", and "High" classes.

- **Strengths:** SVM improves predictions for the "Very High" class with 7166 correct predictions. It also demonstrates marginal improvement over LR in identifying "Low" (16 correct predictions) and "Medium" (6 correct predictions) classes.
- **Weaknesses:** Despite improvements, the SVM model misclassifies the majority of "Medium" and "High" instances as "Very Low", highlighting its difficulty with underrepresented classes.

Long Short-Term Memory

Figure 5.5 shows the strengths and weaknesses of the LSTM model:

- **Strengths:** LSTM achieves comparable results to SVM and LR for the "Very Low" class, with 113,329 correct predictions. It also performs the best among the three models for the "Very High" class, achieving 7246 correct predictions.
- **Weaknesses:** The LSTM model fails to predict any instances of the "Low", "Medium", and "High" classes, misclassifying them all as "Very Low". This suggests significant overfitting to dominant class patterns, likely due to insufficient representation of minority classes in the training data.

Comparative Insights

Across all three models (Figures 5.3–5.5), the significant class imbalance skews predictions heavily towards the "Very Low" and "Very High" categories:

- SVM (Figure 5.4) demonstrates slightly better performance for the "Low" and "Medium" classes compared to LR and LSTM.
- LSTM (Figure 5.5) outperforms both LR and SVM for the "Very High" class but performs poorly for mid-range classes due to overfitting.

These findings highlight the need for techniques to address class imbalance, such as oversampling minority classes or employing advanced cost-sensitive learning methods to improve classification performance for underrepresented categories.

5.6 Discussion

This study investigated the integration of IoT technology with machine learning models to classify stress levels based on environmental factors, specifically temperature and humidity.

The results highlighted the effectiveness of various machine learning approaches, focusing on Logistic Regression (LR), Support Vector Machine (SVM), and Long Short-Term Memory (LSTM) networks. Fuzzy logic was also employed to enhance the classification process, particularly in dealing with the inherent uncertainty and variability in the environmental data. This section discusses the implications of these findings, the limitations of the proposed methodology, and potential avenues for future work.

5.6.1 Analysis of Model Performance

The comparative evaluation of LR, SVM, and LSTM models provided critical insights into their performance. LR and SVM demonstrated high accuracy and precision, making them well-suited for scenarios where computational efficiency and reliability are essential. However, their limitations in capturing temporal dependencies in environmental data were evident.

The LSTM model, which is designed to leverage sequential data, showed moderate performance relative to LR and SVM. While it excelled in identifying temporal trends and fluctuations, its overall accuracy and F1 score were comparatively lower. This discrepancy may stem from the limited size of the training dataset or suboptimal hyperparameter configurations. Nevertheless, the LSTM's capacity to forecast stress-inducing conditions highlights its potential for real-time applications requiring predictive capabilities.

5.6.2 Impact of Environmental Factors

The analysis revealed a strong correlation between environmental factors—temperature, humidity—and stress levels. These findings align with existing research, confirming that elevated temperatures and humidity levels are associated with higher stress. The thresholds identified for temperature and humidity provide actionable insights for designing interventions, particularly in enclosed environments such as offices and homes. These results underscore the importance of addressing environmental conditions as a critical factor in emotional well-being.

5.6.3 Limitations

Several limitations of this study must be acknowledged. Firstly, the dataset was specific to tropical cities, which may limit the generalisability of the findings to other climates and populations. Expanding the dataset to include a broader range of environmental conditions and user demographics could improve the robustness and applicability of the models.

Secondly, while the models effectively classified stress levels, relying solely on temperature and humidity as predictors may oversimplify the complex nature of stress. Incorporating additional variables, such as air quality, noise levels, or physiological data (e.g., heart rate), could provide a more holistic understanding of emotional states.

Lastly, labelling stress levels based on predefined thresholds for temperature and humidity may not account for individual differences in stress responses. Future work should explore personalised models that adapt to users' unique sensitivities to environmental factors, enhancing the system's relevance and accuracy.

5.7 Conclusions

This study demonstrated the potential of integrating IoT and machine learning technologies for stress classification using environmental data. LR and SVM models were shown to be effective for classification tasks, while the LSTM model exhibited promise for predictive applications. Despite certain limitations, the findings lay the groundwork for the development of advanced stress monitoring systems that leverage real-time environmental data. By addressing the identified challenges and exploring proposed future directions, this work has the potential to significantly enhance human well-being in smart environments.

Future Directions

Building on the current findings, several avenues for future research and development are recommended:

- **Dataset Expansion:** Collecting larger datasets from diverse environments and populations will enhance model accuracy and generalisability.
- **Feature Augmentation:** Expanding input features to include physiological signals, activity data, or contextual information will provide a more comprehensive understanding of stress factors.
- **Edge Device Deployment:** Developing lightweight versions of the models for deployment on edge devices can enable real-time stress monitoring and feedback in resource-constrained settings.
- **Adaptive Models:** Designing adaptive models that learn individual users' stress responses over time can improve the system's personalisation and accuracy.

- **Integrated Intervention Systems:** Combining stress monitoring with intervention mechanisms, such as environmental adjustments or relaxation prompts, can create a holistic stress management solution.
- **Data Augmentation and Class Imbalance Mitigation:** Addressing the challenge of class imbalance in stress detection systems, and incorporating techniques such as oversampling, synthetic data generation, or cost-sensitive learning to create more robust models, can improve performance and fairness, especially when dealing with underrepresented stress levels in the data.

These directions aim to refine the system and broaden its applicability, ensuring a more impactful contribution to the field of stress management and mental well-being.

Chapter 6

Prelude to Gamified AI-Driven Depression Monitoring and Management

6.1 Introduction

Chapter 7 examines the behavioural and user-engagement dimension of the proposed IoT-based depression management approach. It presents a gamified AI-driven system that combines therapeutic micro-games, adaptive self-reporting, environmental context, and machine learning analytics and is designed to support depression monitoring and self-regulation. In this chapter, the term *therapeutic micro-games* refers to brief game-based activities designed with therapeutic or well-being-oriented objectives. The role of Chapter 7 in the thesis is to explore how an IoT-based mental health system can move beyond passive sensing and prediction toward active user interaction, reflection, and personalised feedback. This chapter builds on the earlier systematic review, which identified fragmentation in IoT-based depression and emotional well-being systems, and on the contextual chapter, which examined environmental data as a potential interpretive layer. Together, these links position Chapter 7 as the behavioural and interaction layer that connects contextual sensing with the later *IoTMindCare* reference architecture in Chapter 9.

Building on the contextual perspective established in Chapter 5, this chapter focuses on the behavioural and user-engagement aspects of IoT-based depression management. The design of a mental health monitoring system should also consider how users may be encouraged to interact with it, reflect on their emotional state, and receive feedback relevant to their changing circumstances [123]. This is particularly important in depression management, where low motivation, fatigue, cognitive dysfunction, psychomotor slowing, and social withdrawal can reduce a person's capacity or willingness to engage with self-guided digital

tools [124, 125]. Therefore, this chapter shifts the thesis from the contextual sensing layer to the behavioural and interaction layer. It examines how IoT-based mental health monitoring can incorporate more engaging and adaptive forms of interaction that may support reflection and self-regulation, rather than merely functioning as a passive data-collection system.

Digital mental health tools have strong potential to expand access to support, but many systems struggle with long-term engagement, retention, and real-world use [126, 123]. Engagement is not only a usability consideration in this context; it can influence the continuity of data collection, opportunities for timely feedback, and users' ongoing interaction with digital support tools between clinical or informal care interactions [127]. Systematic reviews of digital mental health interventions have emphasised the need for better engagement strategies, personalisation, reminders, feedback, and user-centred design [128, 123].

Gamification has been proposed as one approach that may support motivation and adherence by embedding reflective or well-being-oriented tasks within interactive experiences [129]. At the same time, ecological momentary assessment and adaptive daily surveys can support real-time emotional monitoring while reducing recall bias and improving contextual accuracy [130]. AI-enabled analytics may support these systems by examining patterns across behavioural, emotional, cognitive, and contextual data, with the potential to inform more adaptive feedback and support strategies rather than relying solely on fixed, non-personalised models [131].

Chapter 7 presents a gamified AI-driven system for depression monitoring and management that integrates therapeutic micro-games, adaptive surveys, environmental context, and machine learning analytics within a unified platform. The chapter focuses on how user engagement, self-reporting, contextual information, and AI-based analysis can be combined to explore how they may support depression monitoring, user reflection, and self-regulation, while addressing the need for more integrated and adaptive mental health technologies.

6.2 Role of Chapter 7 in the Thesis

Within the overall thesis, this chapter contributes a user-facing and engagement-oriented design perspective that later informs the integrated architecture. The systematic review identified fragmentation in existing IoT-based depression monitoring systems and motivated a more integrated approach. Chapter 5 explored contextual sensing through environmental data, and Chapter 7 extends this development by exploring how users may interact with the system through game-based tasks and adaptive self-reporting, while behavioural and contextual indicators are collected to inform AI-based analysis.

Chapter 7 is positioned as a bridge between contextual data collection and the later architecture-level contribution of the thesis. It illustrates how the design of mental health monitoring systems can move beyond passive observation by incorporating more interactive, personalised, and adaptive mechanisms. The gamified components are not included solely for entertainment; they serve as low-burden interaction points through which cognitive performance and response patterns can be recorded over time, while mood and engagement-related information can be obtained through self-reporting and interaction data. Similarly, the adaptive survey component supports self-reflection while generating structured emotional data that can be combined with environmental and behavioural inputs.

The main contribution of this chapter is the design and simulation-based evaluation of a gamified AI-driven system that supports both depression monitoring and self-regulation. The chapter contributes to the thesis in four main ways:

1. Demonstrating how gamified micro-interventions can be incorporated into an IoT-based mental health platform, with the aim of encouraging repeated user interaction and capturing behavioural indicators relevant to depression.

2. Showing adaptive survey logic in a simulation-based setting as a potential approach to reducing repetitive prompts while retaining self-reported emotional information. Actual reductions in user burden were not empirically evaluated.

3. Extending the contextual argument developed in the previous chapter by incorporating environmental variables such as temperature and humidity into the behavioural monitoring process.

4. Providing a simulation-based evaluation framework that allows machine learning models and adaptive logic to be tested safely before real-world deployment.

The use of synthetic data is important at this stage of the thesis. Because mental health data is sensitive, early testing with simulated users allows the system logic, model pipeline, and adaptive feedback mechanisms to be explored without exposing real participants to risk. Chapter 7 provides a simulation-based design exploration that can inform future empirical evaluation using real-world behavioural, physiological, and self-reported data.

6.2.1 Synthetic Data Generation, Label Construction, and Evaluation Limitations

Chapter 7 presents a simulation-based evaluation using synthetic data. As described in Section 7.5.2, the dataset comprised 3000 virtual users with 10–30 sessions each. Behavioural,

cognitive, emotional, and environmental features were generated using normal, uniform, and Poisson distributions, probabilistic relationships, and temporal variation. Selected correlations (0.3–0.6) were introduced based on literature-informed assumptions.

As described in Section 7.5.4, depression-related scores were calculated using weighted combinations of the generated features, normalised to [0, 1], and converted into binary labels using the median threshold. The exact numerical weights, their signs, and median tie-handling rules were not fully specified. These labels represent synthetic classifications.

Chapter 7 is an exploratory demonstration of gamified interaction, adaptive surveys, and machine learning integration. Future evaluation requires independent clinical outcomes, participant-disjoint validation, and real-world longitudinal data.

6.3 Connection to Subsequent Chapters

By introducing a gamified, adaptive system, Chapter 7 strengthens the thesis that effective IoT-based depression monitoring requires more than sensors and prediction models. It also requires mechanisms for engagement, self-awareness, personalisation, and contextual interpretation. These ideas directly inform the later architecture in Chapter 9, where the different layers of sensing, analytics, feedback, safety, and stakeholder interaction are brought together into a broader reference architecture.

Chapter 7 presents the design, implementation, and simulation-based evaluation of this gamified AI-driven system. It is an engagement-oriented contribution within the larger thesis, showing how IoT-based mental health monitoring can become more interactive, adaptive, and supportive for users while also providing a research-ready platform for testing future AI models and intervention strategies.

Chapter 7

A Gamified AI-Driven System for Depression Monitoring and Management

7.1 Abstract

Depression affects millions of people worldwide and remains a significant challenge in mental health care. Despite advances in pharmacological and psychotherapeutic treatments, there is a critical need for accessible and engaging tools that help individuals manage their mental health in real time. This paper presents a novel gamified, AI-driven system embedded within Internet of Things (IoT)-enabled environments to address this gap. The proposed platform combines micro-games, adaptive surveys, sensor data, and AI analytics to support personalized and context-aware depression monitoring and self-regulation. Unlike traditional static models, this system continuously tracks behavioral, cognitive, and environmental patterns. This data is then used to deliver timely, tailored interventions. One of its key strengths is a research-ready design that enables real-time simulation, algorithm testing, and hypothesis exploration without relying on large-scale human trials. This makes it easier to study cognitive and emotional trends and improve AI models efficiently. The system is grounded in metacognitive principles. It promotes user engagement and self-awareness through interactive feedback and reflection. Gamification improves the user experience without compromising clinical relevance. We present a unified framework, robust evaluation methods, and insights into scalable mental health solutions. Combining AI, IoT, and gamification, this platform offers a promising new approach for smart, responsive, and data-driven mental health support in modern living environments.

Publication Note

This chapter is based on the published article:

Zamani, S., Rostami, A., Nguyen, M., Sinha, R., & Madanian, S. (2025). A gamified AI-driven system for depression monitoring and management. Applied Sciences, 15(13), 7088.

The published article is available online via https://doi.org/10.3390/app15137088.

The remainder of this chapter is based on the published article, with minor formatting adjustments made to align it with the structure and argument of this thesis.

7.2 Introduction

Depression affects over 280 million people worldwide and is a major global health challenge. It is marked by ongoing sadness, distorted thinking, and a loss of interest or pleasure. Beyond its personal toll, depression has wide-reaching social and economic impacts [4]. Many people still lack acceptable treatment due to barriers like limited access, high costs, and a lack of personalized care. Digital mental health tools offer scalable alternatives but often fail to keep users engaged, adapt to individual needs, or use real-life behavioral and contextual data effectively [132].

To address these limitations, recent work has explored complementary approaches. Gamification enhances engagement and adherence by embedding therapeutic goals into interactive micro-games [129]. Environmental sensing captures contextual variables such as temperature and humidity, which influence mood fluctuations [24]. Ecological momentary assessments (EMAs), implemented via adaptive daily surveys, support real-time self-reflection and minimize recall bias [133]. Furthermore, artificial intelligence (AI) enables dynamic behavior modeling over time, supporting early detection of mental health deterioration and delivering personalized feedback.

The significance of this research lies in its holistic integration of these complementary strategies. While prior studies have examined these components in isolation, few have unified them into a coherent, interactive, and research-ready platform. The core objective of this study is to design, implement, and evaluate an AI-driven framework that synthesizes gamification, adaptive assessment, and environmental sensing to support both individual self-regulation and scalable mental health research.

To guide this investigation, we pose the following research questions:

- RQ1: Can integrating gamified micro-interventions, adaptive ecological surveys, and environmental sensing enhance the early detection of depressive symptom patterns?

- RQ2: How can an AI-enabled digital mental health system dynamically personalize user experiences based on multimodal behavioral and contextual inputs?
- RQ3: How can such a system support end-user self-regulation and researcher-led experimentation within the same operational framework?

To address these questions, we developed a modular digital platform comprising a mobile app, a secure cloud-based infrastructure, and an AI-powered analytics engine. The system collects multimodal data through cognitive games, adaptive daily surveys, and environmental sensors. A synthetic dataset was generated to simulate diverse behavioral patterns, and predictive models were trained using stratified sampling and cross-validation. The platform was evaluated through scenario-based simulations to assess its capability for pattern recognition, risk estimation, and personalized feedback delivery.

This research illustrates how a digital mental health tool can serve therapeutic and research functions by embedding real-time analytics, adaptive user feedback, and simulation tools into a unified system. The platform's design fosters self-awareness and metacognitive regulation, helping users shift from maladaptive thought patterns to healthier strategies while enabling researchers to refine algorithms and hypotheses in real time, iteratively. While synthetic data enables controlled testing of the system's logic and AI models, future studies will involve empirical validation using real-world behavioral and physiological data. This will allow for evaluation under naturalistic conditions and improve clinical applicability.

This work makes three key contributions:

- System design: A novel, AI-driven architecture integrating gamified micro-interventions, adaptive surveys, and environmental data for depression monitoring.
- Simulation and evaluation framework: A synthetic data pipeline for safe, large-scale testing and validation of mental health algorithms.
- Dual functionality: A unified platform serving user self-regulation and researcher experimentation, simultaneously advancing clinical and scientific goals.

This paper is structured as follows: Section 7.3 reviews foundational literature on gamification, EMAs, and AI in mental health. Section 7.4 outlines the system architecture and design rationale. Section 11.5 presents the experimental methodology, and Section 11.6 discusses key challenges and future directions. Through this integrative framework, we aim to inform the development of responsive, scalable, and clinically relevant digital mental health solutions.

7.3 Background and Related Works

This section reviews foundational work on digital mental health tools, focusing on gamification, adaptive daily surveys, AI-based personalization, and the role of environmental features. These themes inform the design and direction of the proposed system and highlight critical areas where existing approaches can be extended or improved.

7.3.1 Gamification in Mental Health Tools

Gamification, applying game design elements such as rewards, feedback, and progression, has gained traction in digital mental health interventions. Systems like SuperBetter and SPARX have demonstrated measurable improvements in resilience and mood, showing that game-like experiences can enhance user motivation and therapeutic outcomes [134, 135].

However, many implementations adopt a fixed design that does not evolve with the user's emotional state, preferences, or behavioral changes over time [129, 136]. This static approach may limit sustained engagement, particularly for users experiencing mood or cognitive functioning fluctuations. Emerging studies advocate for more adaptive frameworks in which gamified elements adjust based on user behavior or input and where AI plays a central role in guiding difficulty, pacing, and feedback [137, 138]. Recent reviews have emphasized the need for personalization in digital interventions to maintain user interest and relevance over time [139].

7.3.2 Adaptive Daily Surveys for Emotional Monitoring

Daily surveys are a widely used tool for emotional check-ins in mobile mental health applications, ranging from structured clinical scales like the Patient Health Questionnaire (PHQ-9) and Generalized Anxiety Disorder scale (GAD-7) to app-specific, simplified self-report instruments. These tools facilitate self-reflection, early symptom recognition, and tracking emotional trends over time [140].

Nonetheless, survey fatigue is a common concern, especially when users are presented with the same set of questions each day. Repetitive and non-contextual prompts can lead to disengagement or incomplete data capture [141]. Researchers have explored adaptive surveys that personalize content based on prior responses, mood history, or contextual factors such as time of day or recent behavior [142, 143]. Adaptive assessments have also shown promise in improving both adherence and the ecological validity of collected data, especially when aligned with users' daily routines [144]. When thoughtfully combined with gamification

elements and feedback mechanisms, adaptive surveys can enhance user experience while preserving data richness and emotional accuracy.

7.3.3 AI for Personalized Depression Support

Artificial intelligence has become essential for identifying depression-related signals from diverse data sources. Techniques such as natural language processing (NLP), supervised learning, and deep neural networks have been applied to detect depressive symptoms from text entries, speech patterns, facial expressions, and mobile usage data [145].

Many approaches rely on passive data collection and generalized models, which may not fully capture individual variations in emotional expression or symptom progression [146]. A growing number of research points toward the benefits of integrating self-report data, behavioral patterns, and environmental signals to create more comprehensive, responsive systems [20]. In particular, in depression, combining multiple data streams and incorporating adaptive learning techniques, such as reinforcement learning, can significantly improve early detection and personalized feedback. Some studies also stress the importance of explainable AI (XAI) to increase user trust and interpretability in clinical contexts [147].

Recent large-scale advancements further support this direction. Zhong et al. [148] introduced a conversational system using large language models (LLMs) that achieved 89% precision in depression detection, outperforming standard clinical tools like the PHQ-9. Similarly, Qiu et al. [149] developed EmoAgent, a multi-agent AI framework designed to ensure emotional safety and stability during human–AI interaction. These systems demonstrate the growing feasibility of real-time, AI-guided mental health monitoring and validate our system's integration of adaptive feedback and explainability.

7.3.4 Environmental Context: Temperature and Mood

Environmental conditions, especially temperature, play a substantial role in shaping emotional states. Elevated temperatures have been linked to psychological distress, irritability, and reduced cognitive performance due to their physiological impact [150, 151]. Conversely, moderate temperatures have been associated with greater emotional stability and improved mood regulation [102].

Seasonal changes, including variations in sunlight and temperature, have also been implicated in depressive conditions such as seasonal affective disorder (SAD), which tends to peak in winter months [152]. Despite this, many digital mental health systems do not consider the environmental context when assessing or responding to user states. Integrating environmental data streams such as weather or ambient light may enable more context-sensitive

interventions and help explain periodic mood shifts [153]. This presents a largely untapped opportunity to enhance the timing and relevance of mental health support, particularly for users sensitive to seasonal or climatic factors.

Complementing these findings, wearable-based sensing has shown promise in capturing environmental and physiological data relevant to mental well-being. For instance, Bloomfield et al. [154] reported that sleep quality indicators captured by the Oura Ring, such as duration, heart rate variability (HRV), and respiratory rate, correlated with perceived stress. Tang et al. [155] developed a deep learning–enabled smart garment that identified sleep stages with 98.6% accuracy, while Fonseka and Woo [156] demonstrated successful clinical applications of wearable pulse-oximetry for home-based mental health assessment. These advances support our inclusion of ambient and sleep-related sensing for enhanced mood–state prediction.

7.3.5 Motivation and Research Direction

These studies establish a solid foundation for digital mental health tools. However, there remains a need for systems that can dynamically adapt to individual users, integrate diverse data types, and balance clinical utility with user engagement. Current tools often operate within narrow design scopes, focusing either on passive sensing, static interventions, or gamification in isolation, without offering a holistic, interactive, and research-ready platform.

Furthermore, while many systems are designed for clinical efficacy, few prioritize user agency, data transparency, or long-term engagement strategies, all critical for sustained use in real-world settings. This research addresses these gaps by proposing an integrated system that combines gamified micro-interventions, adaptive self-reporting, environmental sensing, and AI-driven analytics. The goal is to support users and researchers through a flexible, modular architecture that enhances engagement, personalization, and scientific exploration in depression monitoring and management.

7.4 Materials and Methods: A Gamified AI-Driven System

This section describes the proposed gamified, AI-driven system for depression monitoring and intervention. Developed with React Native 0.78 and Expo for cross-platform access, it leverages AWS for secure, scalable data processing, real-time analytics, and synthetic data generation. The system integrates therapeutic gameplay, adaptive surveys, environmental context, and AI analytics into a unified digital mental health solution.

7.4.1 System Architecture

The system employs a microservices-based, cloud-centric architecture grounded in human-computer interaction and cognitive behavioral design principles [157, 158]. It consists of three core layers:

- Multi-platform application: A cross-compatible user interface supporting Android, iOS, and web platforms to ensure broad accessibility.
- Cloud-based infrastructure: A secure and scalable back end hosted on AWS, utilizing DynamoDB (a NoSQL database) for efficient data storage and management. This infrastructure supports API services, user authentication, and real-time data processing.
- AI analytics and simulation core: This layer leverages AWS SageMaker to implement machine learning models that analyze user behavioral patterns, estimate risk levels related to depressive symptom severity and relapse, and simulate synthetic scenarios for testing intervention strategies. The risk estimation component assesses the likelihood of adverse mental health episodes, enabling timely personalized interventions.

Multi-Platform Application

The application, developed with React Native and Expo, captures comprehensive behavioral data using therapeutic micro-games, mood self-assessments, and adaptive surveys. Location-based services, including weather data, capture environmental context to support affective state modeling. All user interactions, timestamps, game metrics, and emotional feedback are securely transmitted via AWS API Gateway.

Cloud-Based Infrastructure

AWS services [157, 158] underpin the system's back end, enabling efficient data handling and intelligent processing:

- API gateway: Manages secure requests between front-end and back-end.
- Lambda functions: Processes incoming data and triggers downstream workflows.
- DynamoDB: Stores structured user records, including mood logs and environmental readings.
- Cognito: Handles authentication and personalized user access.

- SageMaker: Facilitates model training, simulation, and batch inference.

Figure 7.1 illustrates the system architecture, including the flow from front-end data capture to AI-powered insights.

AI Analytics and Simulation Core

The back-end infrastructure supports advanced analytics through AWS SageMaker and Python-based ML integrations. Data fusion and time-series modeling are employed to identify behavioral trends, predict mood variations, and deliver tailored feedback. AI models are trained to analyze combined indicators, including gameplay performance, survey responses, and environmental variables, offering insights into individual mental health trajectories. Figure 7.2 displays a snapshot of stored user data, capturing gameplay metrics, mood ratings, and contextual information.

7.4.2 System Component

The proposed system includes therapeutic micro-games, adaptive daily surveys, environmental data collection via location services, and AI-driven analytics. Figure 7.3 shows the main application interface, allowing users to select games while weather data is contextualized in real time.

Key components of the system are as follows:

- **Therapeutic micro-games:** The system incorporates lightweight micro-games designed to assess and support cognitive and emotional functioning. These games target attention span, decision-making speed, working memory, and error sensitivity, domains that are frequently affected by depression and stress-related disorders. Prior research has demonstrated that game-based tasks can capture fine-grained behavioral markers and serve as engaging interventions to support emotion regulation and cognitive flexibility [159, 160]. Short, interactive games reduce friction and increase adherence, making them suitable for frequent self-monitoring and low-intensity therapeutic engagement. Figure 7.3 illustrates an example game interface.

- **Adaptive daily surveys:** The platform employs short daily self-report surveys grounded in EMA principles, focusing on core themes such as mood, stress levels, energy, sleep quality, social interaction, and cognitive clarity. Survey questions are context-aware and adapt dynamically based on user history and behavioral data, ensuring relevance while reducing repetition and survey fatigue. This adaptive approach enables high-frequency yet low-burden mental health monitoring, enhancing the granularity and accuracy of symptom tracking [130].

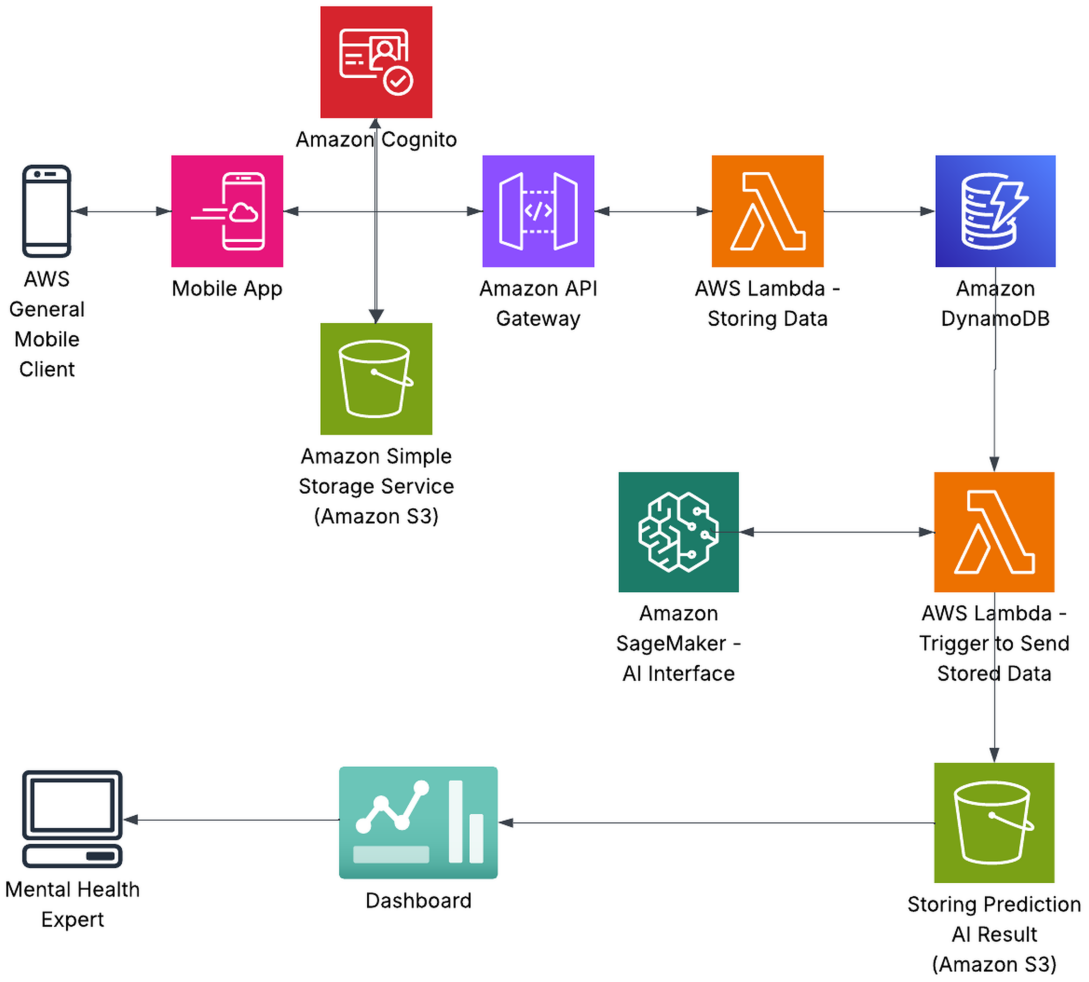


Fig. 7.1 Physical view of the system architecture showing mobile front-end, AWS back-end integration, and AI analytics components.

- **Environmental context awareness:** Real-time environmental data, including temperature, humidity, and light exposure, are collected to provide contextual layers for emotional assessment. This data is correlated with user-reported mood states to explore environmental influences on mental well-being.
- **AI-driven analytics:** The platform leverages machine learning models to integrate data from micro-game performance, adaptive surveys, and environmental sensors. These models first identify hidden behavioral patterns and correlations, such as changes in cognitive speed, mood variability, and environmental influences, indicative of emerging depressive trends. By learning from historical data, the system can predict potential

```
{
    "email": "abc@gmail.com",
    "timeStamp": "2025-02-03T04:37:17+13:00",
    "dayOfWeek": "Mon",
    "game": "Sudoku",
    "mood": 4,
    "longitude": 174.7461,
    "latitude": -36.8667,
    "temperature": 27,
    "cloud-percent": 20,
    "wind": 21.2,
    "city": "Auckland",
    "suburb": "Mt Eden",
    "timeSpent": 21
}
```


Fig. 7.2 Sample of stored user data including mood rating, location, and gameplay interaction timestamps.

mental health risks and deliver real-time, personalized feedback aligned with each user's evolving profile.

7.4.3 Embedded Research Capabilities

The system includes a synthetic data generation module to support robust algorithm validation by simulating user data across various depression states. This enables thorough model evaluation without relying on real participants. Dynamic simulation protocols allow researchers to explore scenarios like rapid mood decline or gradual recovery, offering insights into symptom progression and intervention effects. Interactive dashboards display aggregated insights, such as cognitive–emotional patterns and population trends, aiding hypothesis testing and real-time validation of adaptive interventions.

Collectively, these components create a scalable, adaptable framework that combines AI analytics with a user-centered design to improve depression support and research.

7.4.4 Technical Implementation Details

All modeling, data simulation, and evaluation processes were implemented in Python 3.10. The primary tools used include the following:

- `scikit-learn 1.4` for random forest and logistic regression models.

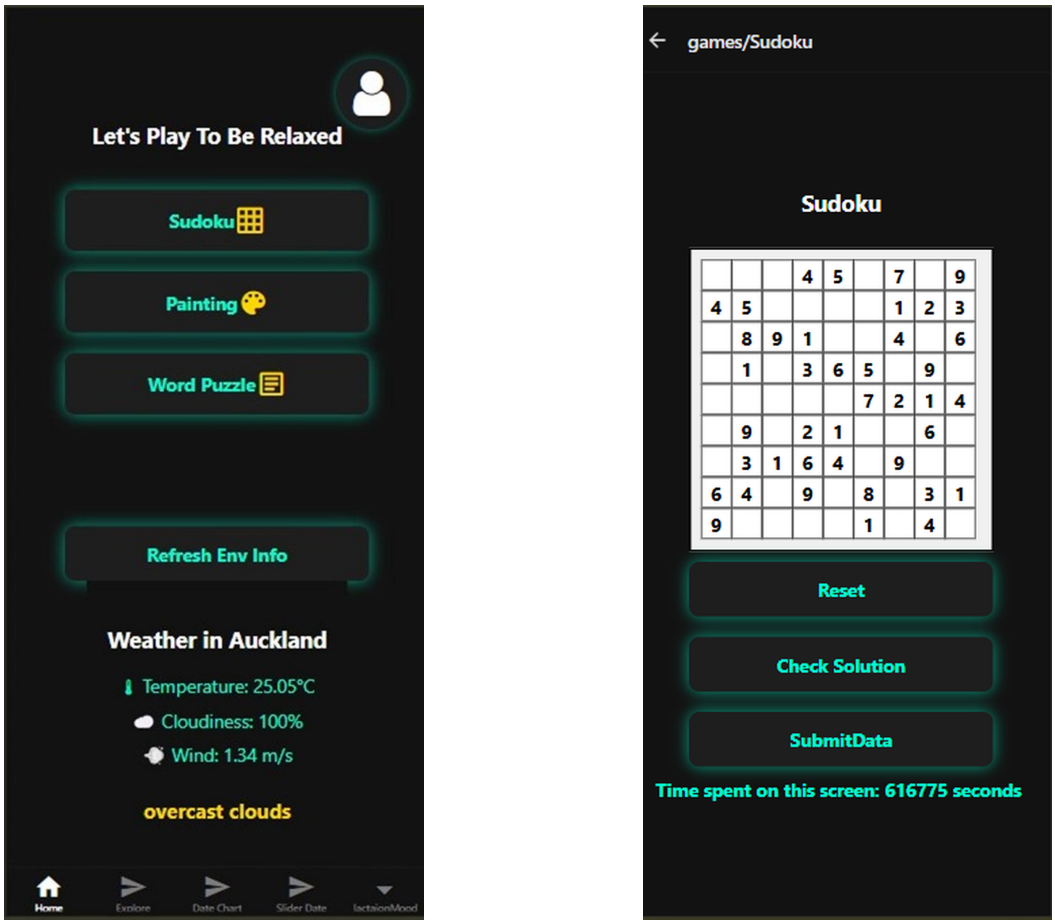


Fig. 7.3 Main application interface showing game selection and real-time weather contextualization, and the game page (Sudoku).

- `TensorFlow 2.15` for the feedforward neural network model.
- `NumPy 1.26`, `Pandas 2.2`, and `Matplotlib 3.8` for data processing and visualization.
- `eli5` for permutation-based feature importance analysis.

7.5 Results: Simulation-Based Evaluation of Depression Detection Models

7.5.1 Evaluation Objectives

This study presents a conceptual evaluation of an AI-augmented system for depression monitoring, using synthetic data to simulate behavioral, cognitive, and contextual variability. The evaluation is structured to assess:

- The capacity of AI models to classify depressive states from multimodal data.
- The feasibility of adaptive survey logic under evolving user states.
- The interpretability and robustness of learned models via visual analytics and feature attribution.

7.5.2 Data Simulation Framework

A synthetic dataset was generated to simulate 3000 virtual users across multiple sessions. Each user profile contains both static (such as baseline cognitive score) and dynamic variables (such as mood score and reaction time). Environmental conditions (such as temperature and humidity) and task metrics (such as accuracy and response latency) were also included, informed by empirical ranges from prior clinical and digital mental health studies [24, 140, 144].

Variables were sampled using Gaussian or uniform distributions depending on empirical variability, with relationships between features (such as sleep and mood, temperature and stress) encoded based on findings from previous literature [150, 102, 153]. Environmental parameters (temperature, humidity) and performance metrics (accuracy, latency) were correlated with vital variables using probabilistic mappings to simulate realistic behavioral patterns.

A continuous *depression-level* score was assigned to each user based on the aggregation of relevant behavioral and environmental indicators. This score was binarized using the median as a threshold to create a balanced binary classification task. Users with scores above the median were labeled as depressed, and those below as healthy. This approach establishes a synthetic ground truth for classification and enables model performance evaluation under controlled conditions.

Synthetic Dataset Structure

The dataset included the following main categories of features:

- Behavioral features: game accuracy, average response latency, error rate, and interaction frequency, representing engagement and performance under cognitive load.
- Cognitive features: baseline cognitive score, working memory load, and attention span indicators were sampled from normal distributions informed by psychological literature.
- Emotional states: self-reported mood, stress level, and energy level, generated dynamically across sessions to simulate fluctuating affective profiles.
- Contextual features: temperature, humidity, and time of day correlated probabilistically with mood and behavioral responses based on findings from environmental psychology studies.
- Temporal dynamics: each synthetic user had 10–30 sessions over time to introduce variability in longitudinal patterns (such as fatigue, recovery trends, or mood swings).

The synthetic data schema emulated real-world variability by combining continuous, ordinal, and binary variables. Gaussian noise and mild correlations (r = 0.3–0.6) were injected between key variables (poor sleep and mood or between temperature and stress) to reflect plausible, ecologically valid relationships observed in prior literature. These features were selected based on their relevance to depressive symptomatology. For example, response latency and game accuracy are linked to psychomotor slowing in depression, while mood variability, stress, and low energy are core affective indicators. Contextual variables like temperature and time of day affect circadian rhythms and mood regulation. Together, these multimodal inputs support realistic simulation of depression-related behavioral and cognitive states. Table 7.1 summarizes the structure of the synthetic dataset.

Table 7.1 Structure of the synthetic dataset categorized by feature, data type, and generation method.

Category	Feature Name	Data Type	Generation Method
Behavioral	Game Accuracy	Continuous	Sampled from a normal distribution
	Response Latency	Continuous	Sampled from normal a distribution
	Error Rate	Continuous	Derived from gameplay
	Interaction Frequency	Count	Poisson distribution
Cognitive	Baseline cognitive score	Continuous	Sampled from a normal distribution
	Working Memory Load	Ordinal	Mapped from the cognitive scale
	Attention Span	Ordinal	Mapped from EMA
Emotional	Mood Score	Continuous	Fluctuates per session
	Stress Level	Ordinal	Dependent on temperature and performance
	Energy Level	Ordinal	Correlated with stress and activity
Contextual	Temperature	Continuous	Realistic seasonal values
	Humidity	Continuous	Realistic seasonal values
	Time of Day	Categorical	Random morning/afternoon/evening split
Temporal	Session Index (1–30)	Integer	Sequential with session-specific noise

7.5.3 Machine Learning Pipeline

Preprocessing

All features were normalized using z-score standardization to ensure comparability across scales. The dataset was split into training (80%) and testing (20%) subsets using stratified sampling to maintain class distribution. Cross-validation was subsequently employed to validate generalizability.

Model Selection and Training

Three supervised classification models were implemented as follows. Each model was trained on the preprocessed feature set, and hyperparameters were manually configured to reduce overfitting while ensuring competitive baseline performance.

- Random forest: An ensemble of 100 decision trees using Gini impurity for feature splitting.
- Logistic regression: A regularized linear model optimized with the LBFGS solver.
- Neural network: A feedforward network with two hidden layers (64 and 32 neurons) using ReLU activation and trained with the Adam optimizer.

Results Summary

Table 7.2 summarizes the classification performance across three machine learning models. The random forest model achieved the highest overall accuracy (0.88) and F1 score (0.878), suggesting it offered the most balanced performance in terms of precision and recall. The neural network model showed the highest ROC-AUC (0.931), indicating strong discriminative ability and potential for detecting subtle depressive patterns. While its F1 score was slightly lower than that of random forest, its sensitivity to class separation could prove valuable in future real-world scenarios.

Table 7.2 Model performance on simulated depression dataset.

Model	Accuracy	Precision	Recall	F1 Score	ROC-AUC
Random Forest	0.880	0.887	0.869	0.878	0.915
Logistic Regression	0.670	0.685	0.616	0.649	0.723
Neural Network	0.865	0.883	0.838	0.860	0.931

While the logistic regression model was less accurate overall, its simplicity and interpretability may be advantageous in clinical or resource-constrained settings. These findings highlight a trade-off between accuracy and interpretability that will be further evaluated in real-world studies.

In addition to classification metrics, our feature importance analysis (Figure 7.4) showed that mood score and environmental features, such as temperature, were the most influential predictors. This confirms the value of multimodal data integration.

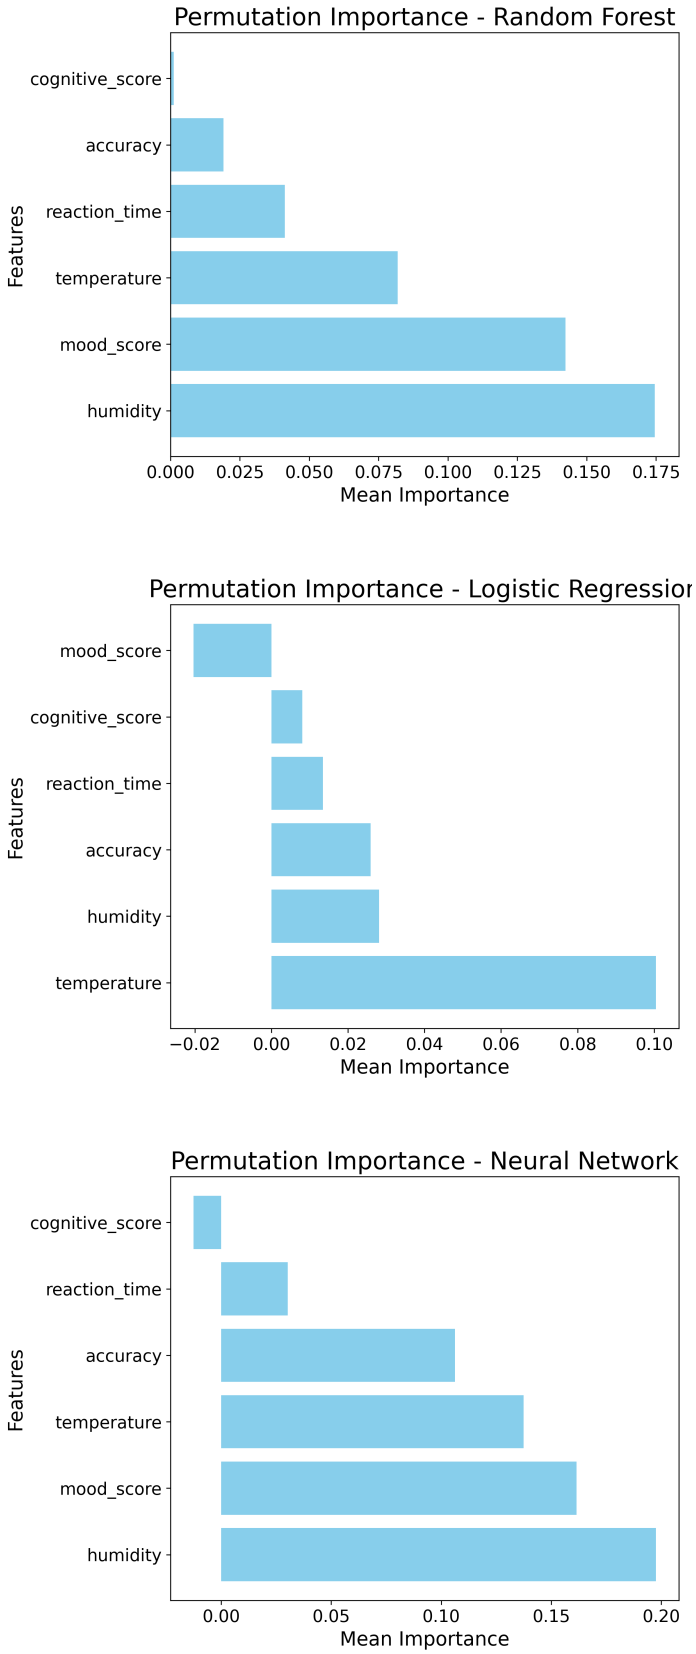


Fig. 7.4 Feature importance based on permutation analysis. Mood score, game accuracy, and temperature are among the top contributors to prediction performance. This highlights the value of integrating behavioral, cognitive, and environmental data.

Furthermore, the adaptive survey simulation over a 12-week virtual period demonstrated the system's capacity to adjust to evolving user states, reducing question fatigue while maintaining data relevance. This validates the theoretical design of the adaptive EMA logic embedded in the platform.

Together, these results establish the proof-of-concept value of the proposed system. The combination of high model performance, real-time personalization, and adaptive feedback highlights the system's potential as both a research platform and a low-intensity self-regulation tool for depression support.

Visualization and Feature Importance

Figure 7.5 presents the ROC curves for all models. The neural network achieved the best area under the curve (AUC), followed closely by random forest. The confusion matrices shown in Figure 7.6 offer additional insight into each classifier's predictive performance and behavior, highlighting their strengths and weaknesses in classifying the data.

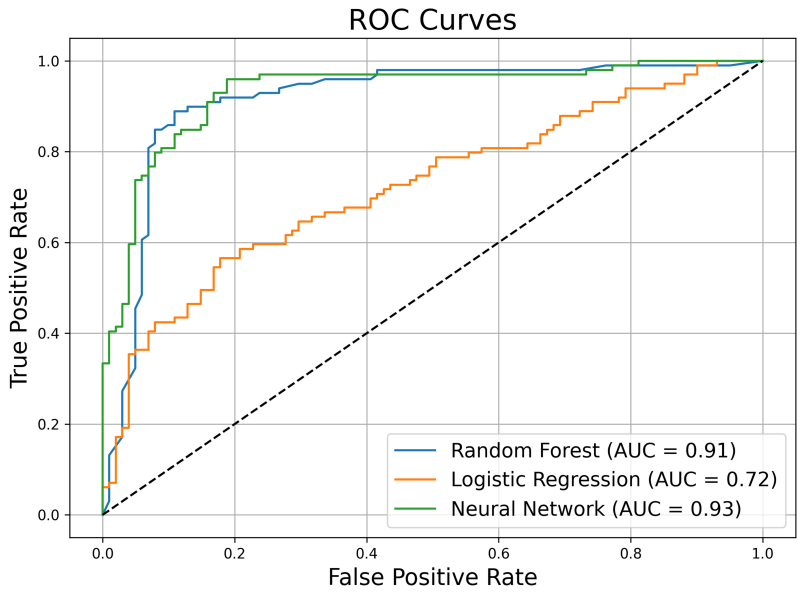


Fig. 7.5 ROC curves for all three models. The neural network shows the highest AUC, indicating strong discriminative performance. Random forest also performs well across all thresholds. Logistic regression has lower sensitivity, especially at lower thresholds.

To assess the contribution of individual features to model predictions, we applied permutation-based feature importance analysis, which measures performance degradation when a feature's values are randomly shuffled. Figure 7.4 shows results across all three models, using the F1 score as the evaluation metric.

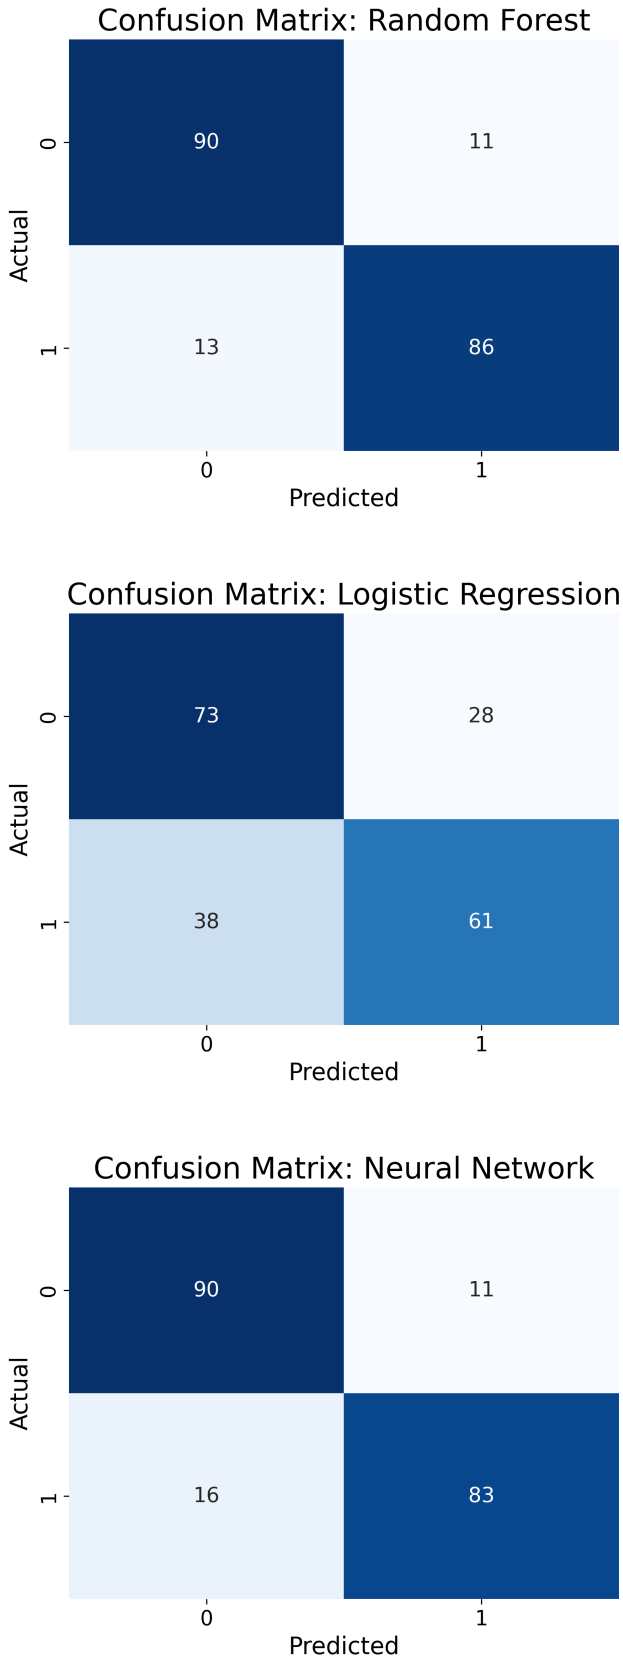


Fig. 7.6 Confusion matrices for each classifier, showing predicted vs. actual labels. Random forest and neural network both achieve high accuracy, with minimal false negatives. Logistic regression shows more misclassifications, particularly in the depressed class.

Across models, *temperature*, *humidity*, and *mood score* emerged as the most influential predictors. These findings align with research linking psychomotor retardation, affective disturbances, and cognitive dysfunction to depression. In contrast, *temperature* and *humidity* had lower importance, indicating a more indirect influence in the simulated context.

7.5.4 Depression Scoring and Classification Algorithm

To assess depressive states, we followed a structured process involving data fusion, composite scoring, and supervised classification:

1. Feature aggregation: Each synthetic user session included behavioral (such as game accuracy and response latency), cognitive (such as baseline cognitive score, attention span), and contextual (such as temperature and humidity) variables.

2. Weighted composite scoring: A continuous depression score D_i was calculated for each user session using the following equation:

$$D_i = \sum_{j=1}^{n} w_j \cdot x_{ij}$$

   where x_{ij} is the value of feature j for session i, and w_j is its empirically assigned weight based on relevance in the literature. For instance, mood variability and energy received higher weights, while environmental features were assigned lower weights.

3. Normalization: Depression scores D_i were min-max scaled to the $[0,1]$ range to standardize the output across sessions.

4. Label assignment: The normalized scores were binarized using the dataset median. Sessions with scores above the median were labeled as “depressed”, while those below were labeled as “non-depressed”. This binarization enabled the construction of a balanced binary classification task.

5. Model training: Supervised machine learning models—including random forest, logistic regression, and neural networks—were trained to predict the binary depression label using the complete feature set. We applied stratified sampling and cross-validation to ensure robustness and generalizability.

7.5.5 Adaptive Survey Logic Simulation

A longitudinal simulation was conducted over 12 virtual weeks to test adaptive survey logic. Survey branching conditions were activated based on trends in fatigue and affective variability. The adaptive system successfully adjusted question content and frequency in response to user state changes, supporting its theoretical responsiveness.

7.5.6 Ethical Considerations

The pilot study involves no real-time data collection or live participants, relying entirely on synthetic data. Ethical considerations ensure that the synthetic data generation process remains unbiased and representative of real-world scenarios. Future real-world data collection deployments will adhere to established ethical standards, including informed consent and data privacy regulations. All analyses were conducted on fully synthetic data. No human subjects were involved. Future deployments will adhere to ethical research standards, including informed consent, data minimization, anonymization, and regulatory compliance, such as GDPR and HIPAA.

7.6 Discussion

The proposed system offers a robust and novel digital mental health monitoring approach by integrating gamification, adaptive EMAs, and AI-driven analytics within an IoT-enabled framework. This section discusses the system's performance, deployment implications, and potential extensions.

The system integrates adaptive self-reporting, gamified micro-interventions, and contextual sensing into a modular, real-time platform. Our simulation results confirm the feasibility of using multimodal data to classify depressive patterns with high accuracy.

7.6.1 Validity of Synthetic Evaluation

Synthetic data enables flexible, early-stage evaluation by facilitating scenario testing, algorithm validation, and stress testing without raising privacy or ethical concerns. Our synthetic dataset was designed to capture realistic behavioral, cognitive, and environmental variability based on empirical ranges reported in prior literature. This approach supports construct validity by simulating plausible depressive states, although it cannot fully replicate the complexity and unpredictability of real-world human behavior.

Model evaluation across three distinct classifiers, random forest, logistic regression, and neural networks, showed consistently strong performance, with the neural network slightly outperforming others (ROC-AUC: 0.931; F1 score: 0.860). These results suggest that the chosen feature space is informative and that depressive patterns are learnable under simulated conditions. We used stratified sampling, cross-validation, and identical training inputs across models to strengthen internal validity and ensure fair comparisons.

To further validate the models' design, the data were simulated using parameter distributions derived from mental health studies. This grounding in empirical findings adds realism to the simulation process and offers confidence in early system performance.

However, generalizability (external validity) remains a key limitation. While the synthetic data reflects diverse behavioral profiles, it may omit latent factors in real populations, such as cultural influences or clinical comorbidities. We aim to address this by validating the platform in upcoming user studies and providing naturalistic data for fine-tuning the models. Ultimately, the synthetic framework enables agile testing of both system logic and adaptive workflows before ethical and logistical challenges of real-world trials arise.

While synthetic evaluation offers a valuable foundation for prototyping and testing, it is not a substitute for longitudinal, real-world validation. These results provide a proof of concept. They are intended to guide future research using real-world data.

7.6.2 Adaptive and Context-Aware Modeling

The system's ability to simulate adaptive surveys over time demonstrates its potential for tracking evolving user states. The survey engine's successful dynamic reconfiguration based on mood and fatigue validates its applicability for long-term mental health monitoring. Additionally, incorporating environmental variables (temperature, humidity) as contextual factors enhances ecological validity, with their impact on model outputs supporting existing research on environmental mood correlations.

This adaptive logic reduces user burden by limiting unnecessary or repetitive prompts and improves the accuracy and personalization of emotional tracking over time. These features position the platform as a sustainable digital mental health support tool, especially in populations prone to disengagement from traditional, rigid app-based assessments.

7.6.3 Research Utility and Generalizability

The system's simulation and dashboard modules turn it into an active research platform, enabling continuous monitoring, performance analysis, and hypothesis testing without repeated real-world trials.

Including visual analytics, such as confusion matrices, ROC curves, and feature importance plots, allows researchers to interpret model behavior and evaluate decision confidence under different scenarios.

However, generalizability is limited by synthetic data, and real-world applications must address issues like missing data, compliance variability, cultural differences, and comorbidities.

A future extension could incorporate transfer learning techniques to adapt the pre-trained models to real-world contexts through transfer learning, improving the practical utility of the system in diverse populations.

Another limitation of this work is the lack of direct benchmarking against existing digital mental health tools or clinical gold standards such as PHQ-9, GAD-7, or commercial platforms like Woebot and Wysa. This was primarily due to the synthetic nature of the dataset and the early-stage focus on system architecture and feasibility. Future work will involve empirical benchmarking using both self-report tools and clinician-validated assessments to compare performance and user outcomes.

At this stage, the objective of the research is to test the hypothesis of having a gamified AI-driven system for depression monitoring and management, and further studies using real data and involving human interaction are required for the system's design and validation.

7.7 Conclusions

This study presents a gamified, AI-driven system for monitoring and managing depression, integrating therapeutic micro-games, adaptive ecological momentary assessments, environmental context sensing, and machine learning analytics. Evaluated using synthetic data, the platform was rigorously tested for classification accuracy, adaptability, and research potential without immediate real-world data collection.

Evaluation results demonstrated high classification accuracy across models, with neural networks achieving an ROC-AUC of 0.931. Feature importance analysis confirmed the relevance of cognitive, emotional, and environmental variables, validating the platform's multimodal design. Neural networks and random forest models outperformed logistic regression across all metrics. While logistic regression offered interpretability, it performed less well on both the F1 score and ROC-AUC. Temperature, humidity, and mood score were identified as the most influential features, consistent with established clinical indicators of depression.

Beyond prediction, the system innovates in adaptive user interaction. Importantly, this system is not intended to replace specialized human intervention but to act as a complementary tool. It offers scalable, real-time support that can bridge gaps in care or assist in self-regulation

between professional sessions, particularly in underserved or low-resource contexts. The longitudinal simulation of survey branching and real-time question adaptation validates the logic engine, which is crucial for sustained user engagement and high-quality data collection. Integrating environmental variables like temperature and humidity adds context-awareness to mental health assessment.

The system's dual functionality, supporting both real-time user feedback and research experimentation, makes it a flexible platform for advancing mental health technologies.

Future Work

The next phase will involve real-world pilot deployments to assess system performance with live data, evaluate long-term user engagement, and improve model generalizability. Integration with passive sensing via wearables and smartphones, including objective monitoring of sleep, activity, and physiological signals, will enhance multimodal data collection. The platform's potential to address other conditions, such as anxiety and stress, will also be explored, advancing it from a conceptual prototype to a scalable clinical tool. This will support more precise, continuous monitoring and improve model robustness in real-world settings.

We plan to conduct empirical benchmarking using established clinical metrics (e.g., PHQ-9 accuracy, engagement retention, symptom tracking sensitivity), comparing our system with tools like Woebot, Wysa, and Moodpath [145].

Chapter 8

Prelude to the IoTMindCare Reference Architecture

8.1 Introduction

Chapter 9 responds to the lack of an integrated architectural foundation for IoT-based depression monitoring and management and marks the thesis's shift toward an architecture-level contribution. Existing IoT-based mental health systems often present useful individual capabilities, such as sensing, mood tracking, self-reporting, anomaly detection, and feedback delivery, as discussed in Chapters 3, 5, and 7. However, these capabilities are frequently developed as separate components rather than as parts of a coherent system. This fragmentation creates challenges for integration, interoperability, safety management, personalisation, stakeholder involvement, and future extension. An integrative reference architecture should organise contextual and behavioural data, adaptive analytics, safety-oriented mechanisms, user-facing feedback, and caregiver or clinician involvement within a coherent system structure.

8.2 Technical Context and Rationale

IoT-based depression management systems present several architectural challenges. First, the system must support multimodal and heterogeneous data sources, including wearable devices, environmental sensors, mobile applications, self-report tools, and behavioural indicators. These data sources differ in availability, reliability, sampling frequency, user burden, and privacy sensitivity. A suitable architecture must accommodate different device configurations and support fallback strategies when particular devices or data streams are unavailable.

Depression monitoring and management require more than data collection and prediction. The system must include mechanisms for interpreting data in context, adapting to individual user baselines, generating feedback, and escalating potential risks when needed. This makes safety and personalisation architectural concerns rather than isolated application features. Safety requires timely detection, appropriate escalation, and reduction of missed high-risk events, while personalisation requires adaptation to individual behaviour, preferences, available devices, and changing circumstances.

The system must also balance local responsiveness with long-term analytics. Edge processing can support privacy, low latency, and continued operation during connectivity interruptions, while cloud processing can support longitudinal analysis, model updates, and broader personalisation. The architecture needs to define how responsibilities are distributed across edge devices, cloud services, and user or caregiver interfaces. These requirements motivate the use of a hybrid edge-cloud structure in the proposed architecture.

8.3 Contribution of Chapter 9

The main contribution of Chapter 9 is the development of *IoTMindCare*, a reference architecture designed to support safe and personalised IoT-based depression monitoring and management. The chapter uses architectural design principles to organise the system around functional requirements, technical constraints, and quality attribute scenarios. Attribute-Driven Design is used to structure the architecture around the identified quality attributes, while trade-off analysis examines how architectural decisions may affect safety, personalisation, latency, privacy, and adaptability. In particular, safety and personalisation are treated as primary quality attributes, reflecting the central concerns of the thesis.

The chapter contributes to the thesis by consolidating lessons from the literature into a structured architectural model. It defines the major layers and components required for IoT-based depression management, including sensing, network and communication, data processing and storage, service logic, and user-facing interfaces. It also introduces safety-related and personalisation mechanisms, including anomaly detection, escalation pathways, feedback integration, user-specific thresholds, and adaptation to heterogeneous device configurations.

In addition, Chapter 9 conducts an architecture-level evaluation of the proposed design through architectural views, quality attribute scenarios, and trade-off analysis. This evaluation shows how different design decisions support or constrain the desired system qualities, particularly safety and personalisation. Chapter 9 provides a design foundation that can

inform future simulation, prototype development, expert refinement, implementation studies, and empirical evaluation.

It is important to distinguish the architecture-level evaluation from empirical validation of system performance. The ATAM-informed analysis examines how the proposed architectural decisions address safety and personalisation requirements and identifies relevant trade-offs. Similarly, the deployment view specifies the intended allocation of components across edge and cloud resources. The findings of Chapter 9 should be interpreted as supporting architectural plausibility and design-level requirements coverage, with operational verification and clinical validation remaining future work.

8.4 Connection to Chapter 11

Chapter 9 establishes the initial integrative architecture for the thesis. However, a reference architecture for depression management must also be examined in relation to expert judgement, stakeholder expectations, safety concerns, and practical feasibility. Chapter 11 builds on *IoTMindCare* through an expert-informed refinement study incorporating technical and clinical perspectives.

This transition is important because the architecture presented in Chapter 9 is developed from literature, architectural reasoning, and system design methods. Chapter 11 strengthens this work by considering how the proposed architecture can be reviewed and refined through expert input. In doing so, the thesis moves from proposing an integrative architecture to examining how that architecture can be improved in light of clinical, technical, ethical, and implementation-related considerations.

Chapter 9

IoTMindCare: An Integrative Reference Architecture for Safe and Personalized IoT-Based Depression Management

9.1 Abstract

Depression affects millions of people worldwide. Traditional management relies heavily on periodic clinical assessments and self-reports, which lack real-time responsiveness and personalization. Despite numerous research prototypes exploring Internet of Things (IoT)-based mental health support, almost none have translated into practical, mainstream solutions. This gap stems from several interrelated challenges, including the absence of robust, flexible, and safe architectural frameworks; the diversity of IoT device ownership; the need for further research across many aspects of technology-based depression support; highly individualized user needs; and ongoing concerns regarding safety and personalization. We aim to develop a reference architecture for IoT-based safe and personalized depression management. We introduce *IoTMindCare*, integrating current best practices while maintaining the flexibility required to incorporate future research and technology innovations. A structured review of contemporary IoT-based solutions for depression management is used to establish their strengths, limitations, and gaps. Then, following the Attribute-Driven Design (ADD) method, we design *IoTMindCare*. The Architecture Trade-off Analysis Method (ATAM) is used to evaluate the proposed reference architecture. The proposed reference architecture features a modular, layered logical view design with cross-layer interactions, incorporating expert input to define system components, data flows, and user requirements. Personalization features, including continuous, context-aware feedback and safety-related mechanisms, were

designed based on the needs of stakeholders, primarily users and caregivers, throughout the system architecture. ATAM evaluation shows that *IoTMindCare* supports safety and personalization significantly better than current designs. This work provides a flexible, safe, and extensible architectural foundation for IoT-based depression management systems, enabling the construction of optimal solutions that integrate the most effective current research and technology while remaining adaptable to future advancements. *IoTMindCare* provides a unifying, aggregation-style reference architecture that consolidates design principles and operational lessons from multiple prior IoT mental-health solutions, enabling these systems to be instantiated, compared, and extended rather than directly competing with any single implementation.

Publication Note

This chapter is based on the published article:

Zamani, S., Sinha, R., Madanian, S., & Nguyen, M. (2025). IoTMindCare: An Integrative Reference Architecture for Safe and Personalized IoT-Based Depression Management. Sensors, 25(22), 6994.

The published article is available online via https://doi.org/10.3390/s25226994.

The remainder of this chapter is based on the published article, with minor formatting adjustments made to align it with the structure and argument of this thesis.

9.2 Introduction

Depression is a prevalent and growing mental health challenge worldwide, affecting over 280 million people [2]. With increasing modernization, rapid urbanization, and the widespread adoption of digital lifestyles, individuals are experiencing heightened levels of stress, loneliness, and social disconnection, factors that contribute to the onset and exacerbation of depressive disorders [161]. Predictive research studies estimate an increase in the number of people affected by depression by 2030 [162]. This trend highlights the urgent need for innovative, accessible, and scalable approaches to support mental health management, especially within environments where individuals lack regular medical oversight.

Traditional depression management methods rely on episodic assessments and manual reporting, which are often prone to recall bias and delays [163]. In contrast, IoT enables continuous and context-aware monitoring via interconnected devices capable of real-time sensing, transmission, and processing. Wearables, ambient sensors, smartphones, and voice assistants can collectively capture multimodal indicators, heart rate variability, sleep cycles,

movement patterns, speech features, and environmental cues to provide a richer, more nuanced monitoring of mental health. This automatic, continuous data stream reduces user burden and has the potential to improve accuracy, enabling a shift from reactive management to proactive and preventive care through immediate anomaly detection and personalized feedback.

Architecturally, most IoT mental health prototypes follow an edge/cloud paradigm [164, 165]: sensors transmit data to a local gateway for initial processing, relaying summaries to cloud platforms for deeper analytics and long-term storage. Fog-computing layers are sometimes inserted to reduce latency during emergencies, and standard protocols such as Message Queuing Telemetry Transport (MQTT) or Constrained Application Protocol (CoAP) provide secure, lightweight connectivity between heterogeneous devices. Despite these advantages, no fully commissioned IoT solution for depression management has been deployed at scale. Our structured literature review, presented in Section 9.3, reveals several persistent barriers, including *integration and interoperability issues*, *privacy and ethical concerns*, *regulatory and compliance challenges*, *technological limitations*, *algorithmic bias*, and *lack of reference standards*.

We hypothesize that a reference architecture that integrates current best practices and prioritizes safety, personalization, and extensibility, can significantly improve the real-world deployment of current and future IoT-based depression management solutions. To test this hypothesis, we follow ADD [166] to develop an integrative reference architecture that integrates the findings from our literature review of current best-practice, including emerging standards and domain requirements, and prioritizes safety and personalization. We then use ATAM [167] to evaluate the resulting architecture. Further validation is performed using Unified Modeling Language (UML) models and comparison with benchmark architectures.

The proposed IoT-based reference architecture (*IoTMindCare*) assimilates best practices, emerging standards, and domain requirements into a unified framework shaped around safety and personalization. It features layering in its logical view to support cross-layer coordination, real-time sensing, behavior analysis, and adaptive interventions through user and caregiver feedback. Unlike existing works, our reference architecture is generalizable and standards-informed. It accommodates heterogeneous sensors, applies to diverse usage contexts, and embeds ethical safeguards from the outset. While clinical trials remain future work, we establish a solid foundation for next-generation mental health platforms that are intelligent, adaptive, and ready for real-world deployment. We present *IoTMindCare* as a collection of established architectural patterns and operational insights that make it easier to adapt existing solutions, compare alternatives, and extend them for future needs. Instead of

emphasizing one specific anomaly detector, the framework provides a flexible architecture for evaluating and selecting components based on application requirements.

Evaluation and comparison show that *IoTMindCare* is designed to support reductions in FPR and FNR through multi-stage verification, feedback-driven adaptation, and edge-cloud tradeoffs. The quantitative evaluation of these reductions is dependent on concrete instantiations and is left to future prototype and pilot studies.

The main contributions of this research are as follows:

- Identifying core challenges in IoT-based depression management via a structured literature review (Section 9.3);
- A modular, extensible reference architecture (Section 9.4) that supports heterogeneous sensors, edge/cloud analytics, safety protocols, and personalization;
- Demonstrating the proposed architecture's applicability by instantiating a concrete IoT depression care system that delivers proactive, ambient, and user-aware support and related evaluations (Section 9.5).

An overview of the proposed *IoTMindCare* reference architecture illustrating its multi-layered structure is indicated in Figure 9.1. The diagram highlights the system's key qualities, including safety, personalization, interoperability, and adaptability.

9.3 Related Works

This review focuses on IoT- and sensor-based solutions that directly inform the design, deployment, and evaluation of reference architectures for continuous,
personalized depression management. We surveyed peer-reviewed publications and technical reports from 2016 to 2025 across IEEE Xplore, ACM Digital Library, PubMed, and Google Scholar using keywords including ‘‘`IoT mental health`’’, ‘‘`depression monitoring architecture`’’, ‘‘`personalized healthcare IoT`’’ and ‘‘`ambient assisted living`’’. We excluded single-purpose apps and non-IoT interventions unless they offered useful contextual integration points. The review prioritizes IoT- and sensor-based solutions that target vulnerable populations, including users at risk of depression and individuals requiring assisted living. This focus ensures that architectural lessons reflect realistic operational constraints and personalization requirements for these user groups.

Prior research clustered into four domains that guided *IoTMindCare*'s architectural drivers and quality-attribute scenarios, as follows:

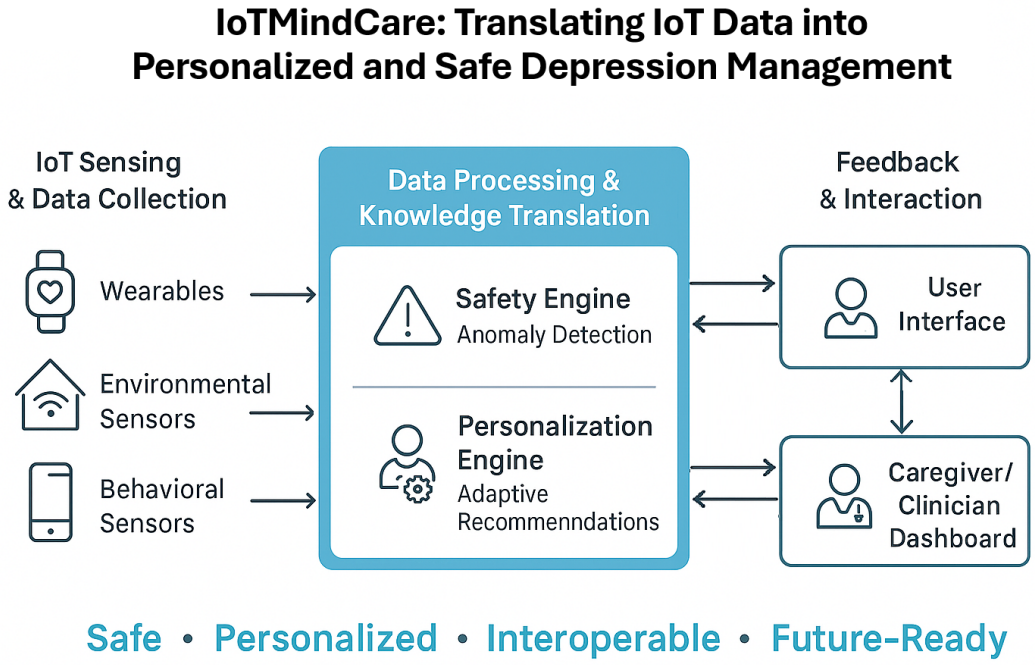


Fig. 9.1 Graphical abstract of the main contributions of *IoTMindCare*

- **Large-scale smart-home and Activities of Daily Living (ADL)/Human Activity Recognition (HAR) platforms:** Large-scale initiatives and datasets, such as the SMART BEAR consortium [168], which targets heterogeneous clinical and consumer sensors for aging-in-place pilots, and contemporary SPHERE [169] releases with multi-sensor annotated datasets, show how diverse sensor suites can be combined for robust ADL inference and longitudinal baselining. Recent work on virtual smart-home simulators and digital twins provides practical tools for testing ADL recognition and failure modes at scale. Meanwhile, contemporary reviews and studies on HAR system lifespan, semi-supervised learning, and ensemble methods highlight modern approaches for handling concept drift and label scarcity in long-term deployments [170]. These efforts collectively provide updated empirical and methodological evidence that multi-modal fusion, simulator-based validation, and hybrid learning pipelines improve robustness and reduce failure modes. These lessons directly motivated *IoTMindCare*'s modular *Sensing Layer*, fallback sensing strategies, and *Personalization Engine* design.
- **Edge–cloud orchestration and privacy-preserving learning:** Edge–fog–cloud orchestration frameworks and federated/on-device learning solutions, such as FedHome

[164], VitalSense [171], and related edge-IoT work, inform placement decisions for training, inference, and data minimization. These works inform our hybrid deployment choices, including edge-first detection for latency and privacy, and cloud-based personalization for long-term adaptation, and shaped the *Personalization Engine* and deployment recipes in *IoTMindCare*. Similarly, Benrimoh et al. [172] explore individualized therapy recommendations using deep neural networks. While its focus is on clinical personalization, *IoTMindCare* complements these approaches by providing an architectural foundation to integrate these models within IoT-based monitoring systems.

- **Safety, escalation, and multi-stage verification:** Ambient Assisted Living (AAL) pilots and alarm-management projects, such as [173] and [174], emphasize the operational need for multi-stage verification and caregiver-in-the-loop escalation to reduce false alarms while ensuring timely response. These operational patterns motivated *IoTMindCare*'s *Safety Engine and tiered alert/escalation policy.*

- **Algorithmic and modality advances (explainability and contactless sensing):** Advances in explainable wearable models and digital-twin/hierarchical detection pipelines provide mechanisms for clinician-facing explainability and hypothesis testing before escalation [175, 165]. Complementary research on contactless modalities, such as radar and Wi-Fi CSI, provides privacy-respecting fallback sensing options when wearables are unavailable [176]. *IoTMindCare* integrates these algorithmic and modality-oriented lessons through explainable interfaces and contactless sensing options, with strict local privacy controls in place.

9.3.1 Operational, Privacy, and Standards

Cross-cutting operational lessons include favoring passive, low-burden sensing for adherence; using multi-modal fusion to reduce false positives; limiting raw-data transmission via on-device preprocessing; and integrating caregiver consent and verification flows.

Standards and middleware, such as HL7 FHIR [177], OpenEHR [178], FIWARE [179], and IEEE 11073 [180], provide concrete data models and connectors that make interoperability and safe clinical integration feasible. *IoTMindCare* adopts these as normative recommendations for semantic integration and secure exchange.

9.3.2 Synthesis: How the Literature Shaped IoTMindCare

The surveyed works are operational inputs rather than competing end-products. They shaped *IoTMindCare* as follows:

- **Sensing layer:** multi-modal fusion, fallback sensing, and passive modalities from ADL/HAR platforms inform sensor selection and local rules.
- **Safety Engine:** multi-stage verification and caregiver-in-the-loop designs from AAL pilots shape alert policies and escalation workflows.
- **Personalization Engine:** federated and edge–cloud learning patterns guide where to place model updates and how to preserve privacy while enabling per-user adaptation.
- **Deployment and standards:** edge–fog–cloud orchestration patterns and interoperability standards guide practical integration and long-term maintainability.

For context, we retain listing therapeutic and self-report systems, such as Woebot [181], Wysa [182], BetterHelp [183], Moodfit [184], and T2 Mood Tracker [185], as possible integration endpoints or manual-input channels.

Table 9.1 summarizes representative IoT, ADL/HAR, and edge-cloud projects and research that directly inform *IoTMindCare* across sensors, personalization, safety, and deployment patterns.

In summary, the literature provides both **Technical Building Blocks**, including sensing, edge–cloud orchestration, safety, and personalization, and **Technical Constraints**, including limited resources, privacy, and device heterogeneity, that collectively informed the architectural drivers, **Quality Attribute Scenario (QAS)**, and component-level decisions in *IoTMindCare*.

9.4 IoTMindCare: Proposed Reference Architecture

IoTMindCare is built using ADD where we first identify key architectural drivers (Section 9.4.1) to the develop architectural views (Section 9.4.2).

9.4.1 Architectural Drivers

Functional Requirements

Our literature review points to five core functional requirements:

Table 9.1 Comparison of recent IoT-based mental and physical health monitoring architectures, their characteristics, and associated benefits.

System / Reference	Focus Area	Key Technologies	Unique Contributions / Features	Limitations Addressed by IoTMindCare
SMART BEAR [168]	Healthy aging, large-scale IoT monitoring	Heterogeneous sensors, AI-based analytics, cloud integration	Personalized interventions for older adults using real-world evidence; scalable data fusion platform	Focuses on physical health; lacks adaptive stress and mood analytics
SPHERE [169]	Multi-sensor home monitoring, ADL/HAR	Video, wearable, and environmental sensing; ML-based ADL inference	Provides open datasets and multimodal monitoring for chronic conditions	Does not address personalization or dynamic emotional context
CASAS [170]	HAR and anomaly detection	Semi-supervised ensemble learning, temporal features	Enhanced HAR accuracy in complex environments	No user-centered personalization or mental health context
VitalSense [171]	Smart city and remote health monitoring	Edge–fog–cloud hierarchy, digital twins, adaptive routing	Low-latency IoT architecture for continuous health monitoring	Lacks domain-specific decision logic for emotional wellbeing
AAL Healthcare 5.0 [173]	Assisted living, healthcare personalization	IoT–AI integration, context awareness, user modeling	Personalized ambient intelligence and health prediction	No explainable AI or caregiver feedback integration
AAL AI Review [174]	AI in AAL	ML and DL algorithms, cognitive computing, privacy concerns	Comprehensive review highlighting challenges in transparency and interoperability	IoTMindCare addresses these gaps via explainable reasoning and standardized APIs
Explainable Wearable Monitoring [175]	Depression/anxiety detection from wearables	Explainable anomaly detection, multimodal fusion	Transparent AI model linking sensor data to mental state changes	Lacks a modular IoT reference architecture or cross-platform interoperability
Contactless IoT Health [176]	Vital signs and ADL tracking	Radar, thermal, and computer vision-based sensing	Comprehensive non-invasive monitoring system with motion tracking	No personalized feedback or multi-stakeholder (user–clinician) data flow
Digital Therapeutic Apps [181–185]	AI chatbots, cognitive behavioral therapy, tele-therapy	Natural Language Processing-based conversational agents, self-tracking apps	Real-time emotional support and mood tracking through mobile interfaces	Standalone, not integrated with physiological or contextual IoT data

F1 **Continuous Monitoring** or sensing of physiological, behavioral, and environmental signals.

F2 **Context-Aware Decision-Making** of multimodal data within spatial-temporal and behavioral contexts.

F3 **Alert Generation and Escalation** to caregivers or clinicians in response to detected risk conditions.

F4 **User and Clinician Interfaces** for stakeholders.

F5 **Personalized Feedback** and intervention mechanisms tailored to individual users.

Quality Attribute Scenarios

Safety and personalization are the key quality attributes in our architecture design process, and are captured using the quality-attribute scenario template [186].
QAS1: Safety: The system must detect critical events, such as depressive crises or self-harm, and respond with timely alerts using a probabilistic, context-aware anomaly detection model (Table 9.2). Detection accuracy, timeliness, and escalation logic are critical. This includes handling uncertainty through a confidence scoring model and minimizing the False Positive Rate (FPR) and False Negative Rate (FNR) as primary operational objectives. Events are cross-referenced with historical and contextual data to improve detection reliability.
QAS2: Personalization: The system must adapt dynamically to users' physiological and behavioral baselines, with support for partial or heterogeneous device configurations (Table 9.3). When sensors are missing, alternative data (manual input, AI-inferred metrics) must substitute for automated sensing. Personalization includes dynamic switching between modes, user-specific thresholds, and real-time updates to decision logic.

Recent IoT-health and applied risk-detection studies explicitly report false positive and false negative outcomes when evaluating anomaly and crisis detection models. For example, Gupta et al. [165] show how hierarchical/federated architectures reduce false positives via multi-layer model aggregation and context validation, motivating multi-stage verification in edge/cloud pipelines. Similarly, applied suicide-risk prediction studies report sensitivity/specificity and precision metrics to quantify clinical detection trade-offs and operational risk [187]. These studies motivate our choice to evaluate safety and personalization using FPR and FNR alongside latency and robustness measures in our reference architecture.

To support reproducible validation of the QASs, we specify measurement protocols and metrics that should be applied when an implementation or simulation of IoTMindCare is constructed. For QAS1 (Safety), we recommend reporting FPR and FNR, sensitivity (recall),

specificity, detection latency (time from event onset to alert), and escalation latency (time from alert to caregiver/clinician notification). For QAS2 (Personalization), we recommend adaptability, installability, customization, co-existence adequacy, and reliability [188]. Validation should use temporally aware splits to reflect longitudinal deployment and expose concept drift. Calibration metrics (e.g., Brier score or reliability diagrams) are recommended when probabilistic confidence scores drive escalation thresholds.

To support reproducible validation of the QASs, we specify measurement protocols and metrics that should be applied when an implementation or simulation of IoTMindCare is constructed. For QAS1 (Safety), we recommend detection metrics such as True Positive Rate / Sensitivity, FPR, FNR, Specificity, and Precision, together with operational metrics, including detection latency (time from event onset to alert) and escalation latency (time from alert to caregiver/clinician notification). For QAS2 (Personalization), we recommend measures of personalization gain, such as absolute and relative reductions in FN and FP rates after per-user adaptation compared to a population model (such as FNR and FPR), model convergence speed (time or number of examples to reach stable per-user performance), and robustness under sensor loss (percent performance drop when a subset of sensors is removed). Validation should use temporally-aware splits (training on earlier windows, testing on later windows) to reflect longitudinal deployment conditions and to expose concept drift. Finally, calibration metrics (such as Brier score or reliability diagrams) should be used where probabilistic confidence scores drive escalation thresholds.

Technical Constraints

The key technical constraints highlighted by our review [164, 165, 189] are:

TC1 **Limited Edge Resources** necessitate light-weight and distributed processing.

TC2 **Privacy Constraints** require local processing and adherence to privacy legislation.

TC3 **Device Heterogeneity** points to the need to tolerate and adapt to variable device configurations.

9.4.2 Architecture Views

The architecture is presented using the 4+1 view model [190]. Each view reports a different but equally critical aspect of the architecture and is linked to a different combination of architectural drivers.

Table 9.2 Scenario for Requirement **Safety**

Stimulus	• The system detects unusual patterns (such as unusual inactivity or negative self-talk) in the incoming data streams.
Stimulus Source	• The user • The system itself • Sensors/devices
Environment	• The system is running under normal operation.
Artifact	• Depression management system components: sensing modules, monitoring modules, communication interfaces, caregiver dashboard, and alert/logging mechanisms.
Response	• The system determines whether observed patterns meet the criteria for a critical or high-risk event and decides whether escalation is required. • If escalation is needed, the system triggers local and remote alerts and records the event for auditing and traceability.
Response Measure	• **Primary objective:** minimize FN (missed legitimate emergencies), aim as close to zero as practicable for the scenario. Studies emphasize maximizing detection sensitivity to avoid missed emergencies, even when this increases false alarms. Zhang et al. [175] reported 12% FN (recall = 0.88), while Wu et al. [164] achieved 95% overall accuracy without explicit FP/FN reporting, providing indicative performance ranges. • **Secondary objective:** minimize FP subject to the primary objective. When FN performance is equal, prefer the lower FP option. In Zhang et al. [175], the corresponding precision is 0.73 implies roughly 27% FP, illustrating the typical trade-off between high sensitivity and increased false alarms. As noted by Gupta et al. [165], multi-stage validation and cross-modal analysis can effectively reduce such false alarms while preserving safety sensitivity. • **Timeliness:** Ensure prompt detection and escalation in a timely manner related to the event. Acceptable latency depends on context, consistent with timelines reported in the SMART BEAR [168] and SPHERE [169] frameworks. • **Auditability:** 100% of the alerts should include contextual and explainable data for human review and post-event analysis, aligning with safety validation requirements observed in prior IoT-based behavioral monitoring systems (Gupta et al. [165] and Zhang et al. [175]).

Logical View

Logically, the architecture has a layered structure due to this pattern's robustness in managing modularity and ensuring logical separation between sensing, processing, and interface components (Table 9.4). Layering is widely adopted in existing solutions [191, 192]. Other

Table 9.3 Scenario for Requirement **Personalization**

Stimulus	• The system needs to tailor its monitoring and intervention strategies to the individual user's characteristics and context (e.g., a user with a specific set of devices and initial profile/context, registers and begins using the system, or an existing user's context changes, new devices, changed environment, or changed clinical needs).
Stimulus Source	• The user (profile, preferences, clinical notes, historical patterns) • The device inventory and environment (available sensors, connectivity) • System-detected context changes (new data streams or changed behavior)
Environment	• Normal operation with diverse user devices, variable connectivity and technology availability, and differing caregiver support models; personalization must operate under these varying constraints.
Artifact	• Personalization Engine and related components in the depression-management system (user profile module, feedback integrator, device-adapter/fallback logic, personalization policy store, and caregiver/clinician interfaces).
Response	• The system can be configured according to the user's identity, context, and available technology, considering who the person is, their mental health needs, and their environmental setup. • The system has high interoperability, adapting to sensor diversity, where multi-device integration supports continuous monitoring across heterogeneous environments. • When specific data streams are missing, the system applies fallback or replacement strategies, consistent with the modular adaptation approaches.
Response Measure	• **Adaptability:** proportion of functions operating under changed conditions (e.g., SMART BEAR adapts to heterogeneous sensors [168]). • **Installability / Customization Rate:** time per configuration/number of supported configurations (e.g., SPHERE supports context-dependent deployments [169]). • **Interoperability Support Rate:** ratio of supported device/format pairs to total expected (e.g., Wu et al. integrate multiple smart home devices [164]). • **Co-existence Adequacy:** ability to work alongside existing systems/devices (e.g., Gupta et al. integrate with existing IoT platforms [165]). • **Replaceability Rate:** percentage of modules swappable without affecting operations (e.g., Zhang et al. propose modular anomaly detection [175]).

architectural styles, such as microkernel and service-oriented architecture, are less suitable

due to weaker support for reactive and real-time event handling required by QAS1, and unacceptable latency for context-sensitive, time-critical operations, respectively.

This view enhances the structure from [193] by embedding user-specific personalization across all layers and enabling dynamic thresholding, offering resilience (TC3) and deeper alignment with QAS2. This enhancement improves system adaptability and responsiveness beyond prior work.

Each layer supports specific architectural drivers. The Sensing layer provides real-time signal capture support (F1, QAS1). The Network and Communication layer ensures reliable data routing (F2, TC3). Data Processing and Storage enables edge/cloud processing (F2, QAS1, QAS2). The Service layer implements core decision logic (F2, F3, QAS1). The Interface layer provides personalized feedback and real-time alerting (F3-F5, QAS2).

Table 9.4 Logical view - layering diagram showing the five architectural layers of the architecture and their mapping to safety and personalization

Interface Layer (user interface on mobile devices and actuators, caregiver's access via dashboard, and feedback mechanism to upgrade the system)	**Personalization** (supporting device variability, customized treatments, user-specific thresholds, and behavioral diversity)	**Safety** (anomalous behavior detection by continous monitoring, real-time alert to caregivers, and environemtal adjustments)
Service Layer (real-time monitoring, automated alerts by decision-making, environmental adjustments)		
Data Processing and Storage (local processing for real-time and rule-based triggers, cloud processing for long-term behavioral analytics, and data storage)		
Network and Communication (local communication, gateway, remote communication, ensuring reliable data exchange between system components		
Sensing (collect multimodal data from physiological, environmental, and behavioral sensors, depending on sensors' availability)		

Development View

The development view (Figure 9.2) defines the internal structure and interactions between software modules distributed across smart home environments, cloud infrastructure, and external actors such as users and caregivers.

The development view adopts a component-based architecture with RESTful and message-passing interfaces due to support of loose coupling, deployment flexibility, and fault isolation [194–196]. This pattern enables modular, distributed services that feature real-time responsiveness and adaptability.

Unlike fixed-rule systems like CarePredict [197] using predefined activity models, *IoTMindCare* dynamically adapts to each user's data through the cloud-based `Personalization Engine`. Feedback loops are integrated via the `Feedback Integrator`, improving the system over time. This enhancement is central to achieving `QAS2` under real-world device constraints (`TC3`), highlighting adaptability and scalability.

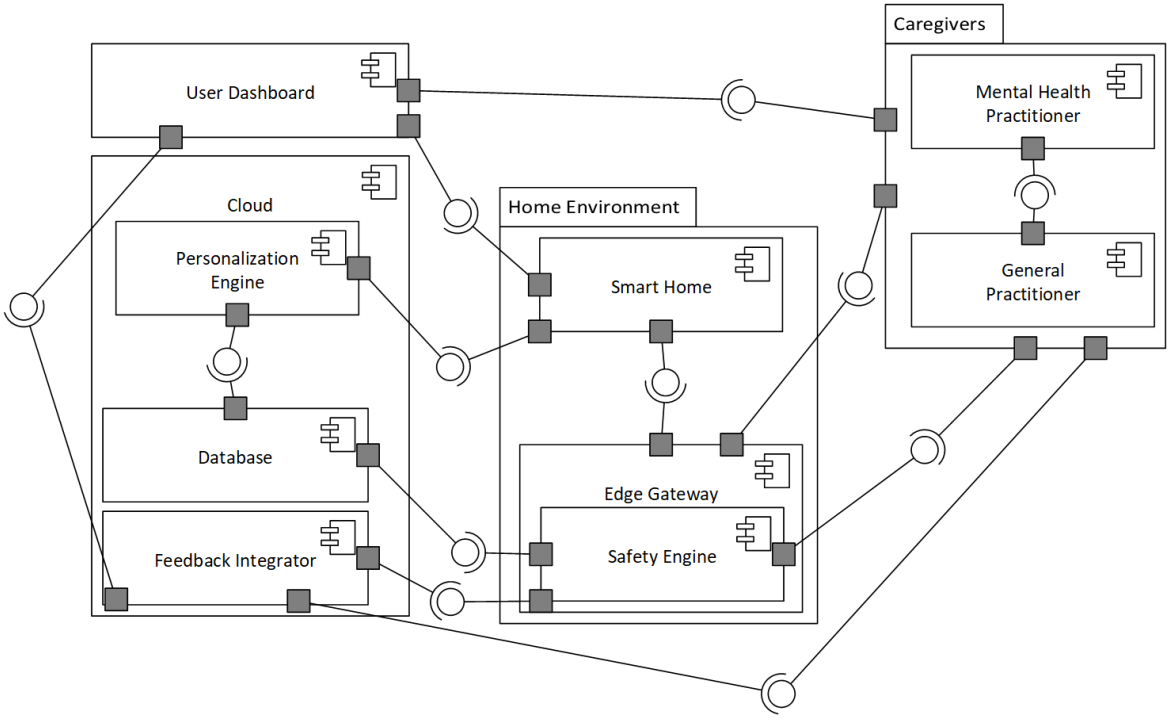


Fig. 9.2 Development view - UML Component diagram showing the major software modules for smart home, cloud, and user environments

`Smart Home` edge interface performs continuous monitoring (`F1`) and forwards data to the `Edge Gateway`, which executes lightweight inference and handles fallback logic during cloud disconnection (`QAS1`, `TC1`, `TC2`). The `Safety Engine` performs anomaly detection using probabilistic models and generates alerts based on severity thresholds (`F3`, `QAS1`). The

cloud-based `Personalization Engine` manages user-specific models and dynamically adjusts service logic based on behavioral patterns (F5, QAS2, TC3). Moreover, the `Feedback Integrator` incorporates feedback from users and caregivers to refine system behavior (F5, QAS2), while the `Central Database` stores longitudinal data and personalization parameters.

For user interaction, `Mobile Application` enables real-time feedback, notifications, and manual data entry for user reflection, alert acknowledgment, and configuration (F4, F5). `Clinician/Caregiver Dashboard` provides real-time and retrospective views of user states, anomaly logs, and behavioral trends for escalation actions and threshold customization (F3, F4, QAS1).

Process View

The process view (Figure 9.3) describes the system's dynamic behavior, capturing the interaction patterns among core components during real-time operation. It adopts an event-driven architectural style, well-suited for reactive, asynchronous IoT-based health monitoring systems due to its responsiveness and scalability [198, 199]. It also uses a client–server pattern, where cloud services hold stateful logic, while edge and user devices act as clients, emitting or receiving events.

Unlike rule-based detection chains in [193], this view incorporates closed-loop adaptation, probabilistic scoring, and scalable escalation, offering resilience and evolving personalization aligned with real-world deployment challenges.

We propose four key processes reflecting the system's end-to-end lifecycle: **System Initialization and Personalization, Continuous Monitoring, Candidate Anomaly Detection**, and **Feedback Integration**. **System Initialization and Personalization** begins during setup or after reconfiguration. `General Practitioner` provides user-specific information such as baseline activity and risk thresholds. The cloud-based `Personalization Engine` generates contextual thresholds and fallback logic for sensor gaps (F5, QAS2, TC3).

In **Continuous Monitoring**, `Edge Gateway` gathers and preprocesses multimodal data. Obvious anomalies are flagged locally, while filtered data is sent to the cloud for deeper analysis. Real-time and historical data are stored centrally (F1, F2, TC1, TC2). In **Candidate Anomaly Detection and Alert Escalation**, `Safety Engine` evaluates deviations using probabilistic models. If an anomaly is detected, multi-level alerts are triggered: from users to general and mental health practitioners, depending on severity (F2, F3, QAS1). Finally, in **Feedback Integration and Adaptation**, users and clinicians provide feedback through dashboards. `Feedback Integrator` refines the personalization logic and detection thresholds, enabling continuous learning (F5, QAS2, TC3).

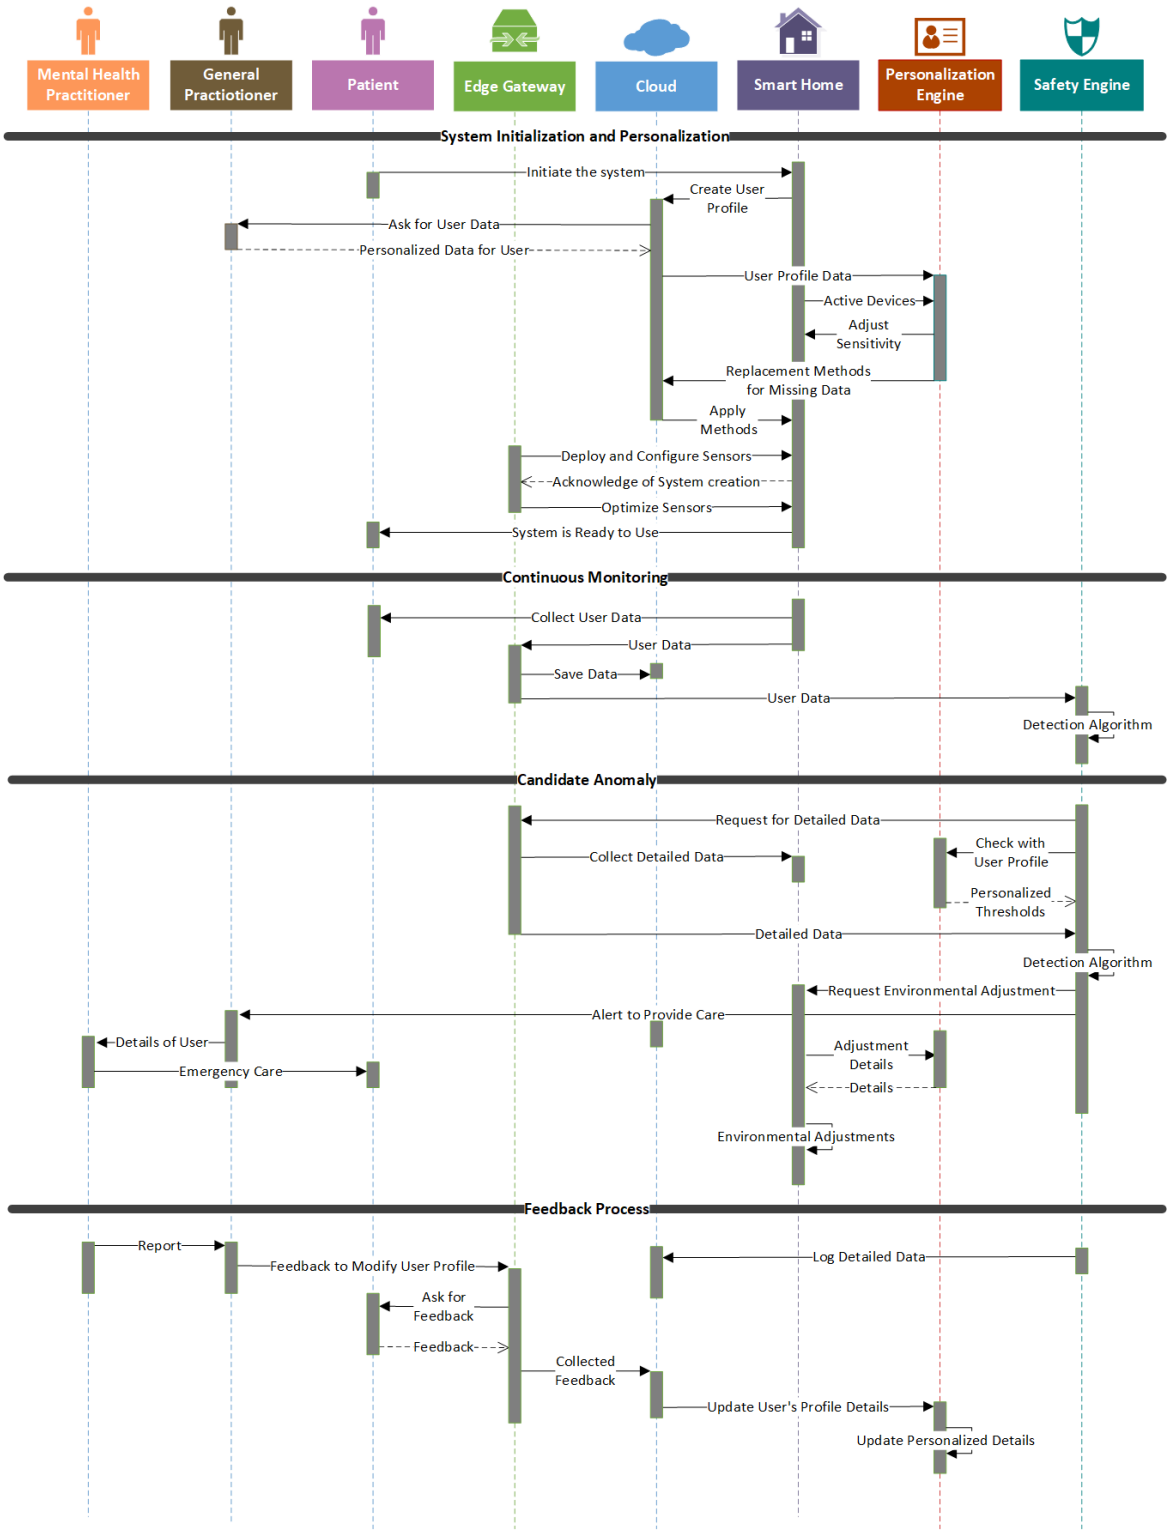


Fig. 9.3 Process view - UML Sequence Diagram showing end-to-end system behavior categorized into four stages and featuring dynamic interactions among smart home, cloud, and user interfaces

Deployment View

The deployment view illustrates how software components are physically distributed across edge devices, cloud infrastructure, and stakeholder interfaces (Figure 9.4). This view relates to latency, fault tolerance, and scalability requirements. A hybrid edge–cloud deployment strategy is adopted, a common best practice in current works [200, 201].

The deployment view enables scalability and resilience by decoupling edge and cloud operations. Cloud components can be updated dynamically without disrupting critical edge functions. Communication protocols are optimized for IoT reliability and low-latency demands, aligning with system constraints (`TC1`–`TC3`) and quality attributes (`QAS1`, `QAS2`).

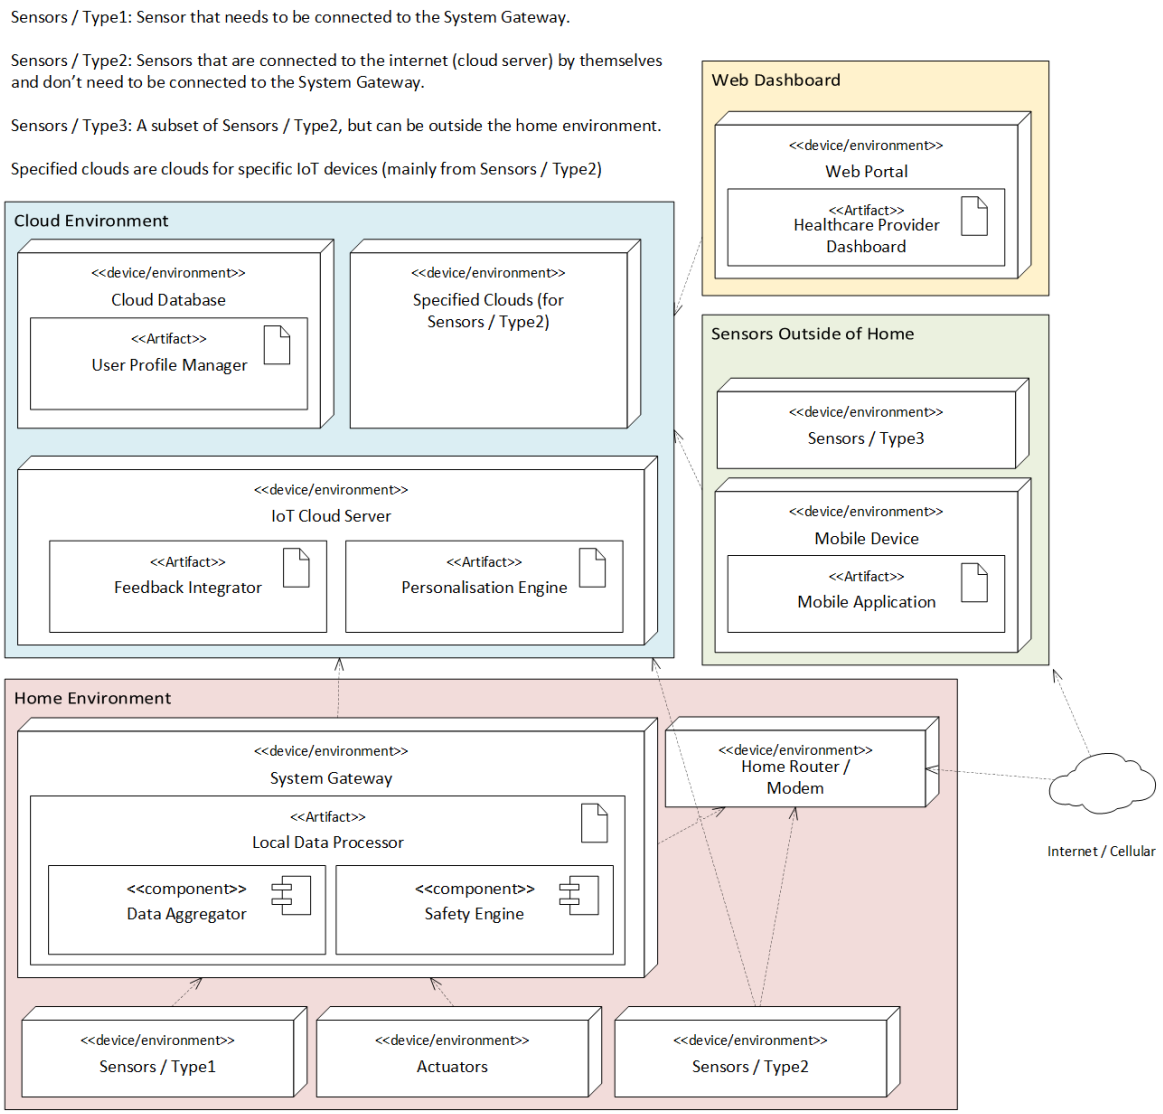


Fig. 9.4 Deployment view - UML Deployment diagram showing the mapping of the development view components to physical or infrastructure resources

Multiple `Sensing Devices` capture multimodal behavioral and physiological data in real time. These feed into a local `Smart Home Gateway`, which hosts the `Edge Gateway` and `Safety Engine`. This setup supports on-site anomaly detection, buffering, and fallback logic during cloud disconnections (`QAS1`, `TC1`).

The `Cloud Infrastructure` manages the `Personalization Engine`, `Feedback Integrator`, and long-term `Database`. It handles high-complexity tasks such as behavioral analytics, adaptive model updates, and secure storage. Encrypted communication between the cloud and smart home ensures privacy and integrity (`QAS2`).

Users and caregivers access the system through `Mobile Applications` and `Clinician Dashboards`, which interface primarily with cloud services for alerts, feedback, and system configuration. These apps can also connect locally for real-time updates when needed.

9.5 Architecture Evaluation

We evaluate the proposed *IoTMindCare* using ATAM [202, 167] for its support for the primary QAs, safety and personalization, as well as complementary 4+1 UML views (See Section 9.4). The QA scenarios from Section 9.4.1 were analyzed to ensure that they are holistically addressed by the architecture, across all relevant views. Table 9.5 details how each QA is supported by *IoTMindCare*, and indicates to a robust architecture delivering safe and personalized IoT-based depression management.

Tradeoff Analysis

The architectural choices must also balance the tradeoffs between *safety* and *personalization*, particularly in minimizing the FPR and FNR in behavioral anomaly detection. We identified the following critical tradeoffs:

1. **Safety vs. Personalisation:** A central tradeoff exists between safety and personalization. Personalization improves detection sensitivity (reducing FNR) by adapting to individual baselines; however, overly narrow per-user thresholds may increase FPR if not validated. IoTMindCare mitigates this via a tiered verification pipeline (edge trigger $\rightarrow$ cloud contextual validation $\rightarrow$ caregiver confirmation) and feedback-driven recalibration. However, overly personalized thresholds risk increasing FPs if not rigorously validated. This tradeoff is mitigated by delegating final anomaly verification to a centralized cloud component (Figure 9.2) and integrating feedback from caregivers (Figure 9.3) to calibrate and improve models iteratively.

Table 9.5 Mapping Quality Attribute Scenarios to architectural evidence

QAS	**Relevant Components / Diagrams**	**Validation Evidence**	**Component Level Mapping (Figure 9.2)**	**Behavioral-Level Mapping (Figure 9.3)**	**Deployment-Level Mapping (Figure 9.4)**
Safety (FPR) (Table 9.2)	`Safety Engine` (Figure 9.4 and 9.2) - `Alert Generation` (Candidate Anomaly Section in Figure 9.3))	Alert verification by checking with the edge gateway to reduce incorrect alerts	Multi-layer verification across the `System Gateway` and `Cloud`	Reduced by context-aware confirmation in the cloud before triggering alerts	Ensuring local anomaly detection continues during network downtime (edge computing on the `System Gateway`)
Safety (FNR) (Table 9.2)	`Safety Engine` (Figure 9.4 and 9.2) - Caregiver and user feedback by `Feedback Integrator` (Feedback Process Section in Figure 9.3)	Feedback integration to detect missed anomalies	`Caregivers` and `User` interface to facilitates feedback loops and missed-event reporting	Addressed via caregiver and user feedback paths feeding into the anomaly model	Ensuring local anomaly detection continues during network downtime (edge computing on the `System Gateway`)
Personalization (Figure 9.3)	`Personalisation Engine` (Figure 9.2 and 9.4) - `User Dashboard` (Figure 9.2 and 9.4) - `User Profile` on `Cloud` (Figure 9.4) - Integrating different sensors (Figure 9.4) - Ask user data from `General Practitioner` 9.3) - Adjust devices and sensitivity (Figure 9.3)	Adaptable thresholds - Behavioral baselining in algorithms - User-specific anomaly models by checking with user profile - Setting up based on the user's devices	`User` module to build user-specific behavioral models and adapt thresholds over time - System configuration based on the user's specified devices	Individualized thresholds used during verification, and using the user's available set of devices	Hosting behavioral models and user data in the cloud enables modifications as user thresholds change

2. **Feedback Loops vs. Timeliness:** Incorporating caregiver and user feedback into anomaly detection loops (Figure 9.3) helps reduce FPR/FNR rates over time. However, this introduces delays in model retraining and potential inconsistencies during transition phases. The architecture mitigates this by decoupling feedback ingestion from real-time alerting, maintaining responsiveness while still enabling long-term model refinement.

3. **Local Autonomy vs. Central Oversight:** To ensure availability and responsiveness, the system gateway performs initial anomaly evaluation at the edge (Figure 9.4). This supports low-latency decision-making and offers resilience during cloud disconnection, contributing to safety (FNR reduction). However, relying solely on edge intelligence limits access to broader contextual information needed to reduce FPRs. As a result, the architecture introduces a tiered decision-making pipeline: edge components perform preliminary assessments, while cloud services provide contextual validation to reduce FPRs. This design reflects a deliberate trade-off between local autonomy and central oversight.

IoTMindCare intentionally aggregates capabilities from the surveyed literature so that many existing architectures can be expressed as subsets or specific instantiations of the reference model. This framing avoids claiming direct superiority; instead, the contribution is a unifying architecture that (a) makes feature/metric tradeoffs explicit, (b) clarifies where variability points should be placed to compare detectors (e.g., choice of anomaly model, thresholding strategy, retraining cadence), and (c) provides concrete recipes for extending prior designs toward more robust, future-proof deployments.

Overall, *IoTMindCare* carefully balances quality attributes, with safety and personalization as primary drivers. Modularity, layered processing, and feedback integration strategies explicitly addressed tradeoffs, as evidenced in the UML diagrams and quality scenario mappings.

9.5.1 Comparison with Existing Architectures

Table 9.8 provides a high-level comparison between IoTMindCare and selected contemporary IoT-health frameworks across data sources, personalization, edge/cloud strategy, explainability, and multi-stage anomaly handling.

This comparison focuses on architectural patterns, operational principles, and design choices rather than empirical performance metrics, as IoTMindCare is a reference architecture.

Table 9.6 Comparison with recent IoT-health architectures

Feature	**Zhang et al.** [175]	**Wu et al.** [164]	**Gupta et al.** [165]	**IoTMindCare**
Data Sources (See `Logical View` in Section 9.4.2)	Wearables	Wearables/ Home Sensors	Wearables/ Digital Twins	Wearables/IoT /Environment
Personalization (See `Process View` in Section 9.4.2)	Global Model	User-Specific Models	Group-Based Models	User Profile/ Feedback-Driven
Explainability (See `Development View` in Section 9.4.2)	SHAP	–	–	Human-in-Loop / Logic Layers and Modules
FPR/FNR Strategy (See `Process View` in Section 9.4.2	Global Threshold	Personalized FL	Multilevel Validation	Multi-Stage / Feedback
Deployment Topology (See `Deployment View` in Section 9.4.2)	-	Edge-Cloud Based FL	Hierarchical Edge–Cloud	Hybrid Edge–Cloud with Fallback and Local Safety Logic

Our architecture combines explainability and personalization within a hierarchical edge-cloud structure, while having other levels of safety and personalization. It extends these through the following:

1. Incorporating home environmental sensors.
2. Employing multi-stage anomaly verification to manage FPR/FNR tradeoffs.
3. Supporting feedback-driven model adaptation with user/caregiver input.
4. Enabling edge-first inference policies for low latency and privacy.

While Zhang et al.'s model [175] presents strong anomaly detection within a controlled dataset, it has a wearable-only scope and centralized deployment, which limits generalization and real-world responsiveness. In contrast, our system is designed for continuous deployment, with added robustness mechanisms (like model updates from feedback) and explainable alerts through transparent processing stages. Compared to federated learning systems, our architecture relies on explicit feedback loops rather than fully distributed updates, offering a more straightforward path to per-user tuning and safety in mental health contexts.

These comparisons suggest that our architecture is a more holistic solution for real-world, mental health IoT deployment, better suited to privacy, safety, personalization, and real-time operation.

9.5.2 Scenario-Driven Instantiation and Configuration

To demonstrate the applicability of *IoTMindCare* in realistic deployments while preserving its role as a high-level reference architecture, we provide a set of concrete IoT-driven scenarios drawn from the literature, along with a mapping of these scenarios to the architecture. For each scenario, we indicate relevant literature sources, the architecture views most involved, the primary QAS affected, and brief configuration guidance. These examples are meant as practical guides for building, testing, or running pilot versions of systems.

Representative Scenarios

1. **Early deterioration detection (gradual mood decline):** Long-term ADL/sleep/activity drift that signals worsening depression (motivated by SMART BEAR [168]).
2. **Acute crisis detection (self-harm/suicidal risk):** Rapid onset of high-risk states requiring low-latency detection and escalation (inspired by AAL pilots and safety studies [174, 173]).
3. **Sleep disturbance and insomnia monitoring:** Night-time patterns detected via bed/pressure and respiration (contactless modalities), indicating sleep disruption associated with mood change (motivated by SPHERE, contactless sensing work [169, 176]).
4. **Social isolation/activity withdrawal:** Reduced outings, fewer device interactions, reduced social communications extracted from ADL/HAR traces (CASAS [203] and SPHERE [169]).
5. **Medication non-adherence/routine deviation:** Missed medication or routine changes inferred from smart-plug, cabinet-door, and activity sensors (SMART BEAR [168] AAL studies [174, 173]).
6. **Acute stress episode (physiological spike):** Short-term HRV and activity spikes detected by wearables requiring brief interventions or prompting (wearable explainability studies [175]).
7. **Contactless monitoring for privacy-sensitive users:** Use radar/Wi-Fi channel state information and on-device inference where wearables are infeasible (contactless systems [176]).
8. **Teletherapy-triggered sensing:** Trigger richer sensing or clinician contact when conversational agents or self-report apps indicate increased risk (integration with digital therapeutics).

Table 9.7 Scenario to architecture mapping: signals, literature sources, relevant views, primary QAS and practical configuration notes for prototype instantiation.

Scenario	Key signals / sensors	Source (examples)	Relevant views	Primary QAS & config notes
Early deterioration detection	Activity patterns, sleep trends, phone/PC usage, longitudinal HR/actigraphy	SMART BEAR [168]	Logical, Deployment, Process	**QAS2 (Personalization).** Longitudinal baselining in Personalization Engine; cloud-hosted model training with scheduled recalibration; edge summarization to preserve privacy. Use drift-detection and periodic retraining hooks; measure FPR, FNR, sensitivity, specificity, detection latency, and convergence rate.
Acute crisis detection	HRV, speech content/prosody, sudden inactivity, environmental context	AAL reviews / safety pilots [174, 173]	Service, Process, Deployment	**QAS1 (Safety).** Edge-first low-latency detection (Safety Engine) with multi-stage confirmation: edge trigger → cloud contextual validation → caregiver escalation. Configure conservative thresholds + high-confidence escalation; log detection latency and sensitivity/specificity.
Sleep disturbance monitoring	Bed-pressure, respiration (radar/CSI), actigraphy	SPHERE [169], Li et al. [176]	Logical, Deployment	**QAS2 (Personalization)** and **QAS1 (Safety).** Use contactless modalities where wearables unavailable; process on edge for privacy; personalize sleep thresholds per-user; validate with sleep-specific sensitivity score and detection-latency metrics.
Social isolation / withdrawal	Reduced movement, decreased device interactions, fewer outgoing calls	CASAS / HAR literature [170, 169]	Logical, Process, Development	**QAS2.** Personalization Engine tracks social-activity baselines; thresholds adapted via Feedback Integrator; A/B style trials to compare static vs adaptive thresholds; measure reduction in FNR for isolation detection.
Medication / routine deviation	Smart-plug, cabinet/door contacts, schedule adherence logs	SMART BEAR [168], AAL reviews [174]	Development, Deployment	**QAS1 & QAS2.** Edge acknowledgement prompts (in-situ) before escalation; local rule filters + cloud validation; integrate clinician-configured escalation policies. Monitor FPR rate of missed-medication alerts.
Acute stress episode	Wearable HR/HRV, accelerometer, short-term speech features	Explainable wearable studies [175]	Logical, Service	**QAS1 & QAS2.** Edge/local detection for low latency, short-term buffering to cloud for contextual validation; use explainability output for clinician dashboards; evaluate with precision/recall and detection latency.
Contactless privacy-sensitive monitoring	Radar, Wi-Fi CSI, on-device feature extraction	Li et al. [176], contactless sensing literature	Logical, Deployment	**QAS2 (and privacy constraints).** Prefer on-device feature extraction and encrypted summaries; disable raw audio/video transmission; define privacy-preserving model update policies (differentially private updates or federated patterns).
Teletherapy-triggered sensing	Self-report, chatbot flags, scheduled clinician prompts	Digital therapeutics / chatbots [181, 182]	Interface, Process, Development	**QAS2.** Use chatbot signals to change sensing mode (e.g., increase sampling), start short-term local analytics, and route data to clinician dashboards; measure user acceptability and engagement.

For a given deployment, practitioners can use the recipes above to configure IoTMindCare as follows:

- **Select sensors and signals** relevant to the scenario
- **Choose model placement:** edge-first for low-latency safety scenarios; cloud-enabled personalization for long-term baselining and retraining.
- **Define escalation policy:** map detection confidence levels to tiered actions (edge prompt, cloud validation, caregiver/clinician alert) within the Safety Engine.
- **Personalize and test:** initialize thresholds from population priors, then adapt per-user via the Feedback Integrator. Use concrete quantitative validation metrics, such as reductions in FNR, reductions in FPR, improvements in sensitivity/specificity, and changes in detection and escalation latency, to evaluate the personalization gain.
- **Privacy and failover:** apply on-device preprocessing, encrypted model updates, or federated learning, and contactless fallbacks where wearables are refused.

These scenario instantiations are explicitly intended to guide future prototype or simulation studies. For each scenario, we recommend a pilot that measures the QA metrics listed above (FPR/FNR, detection latency, and personalization gain), with clearly defined A/B comparisons or pre- and post-personalization analyses. These evaluation plans are consistent with the ATAM-driven analysis and the QA metric framework introduced in Section 9.4.1.
Empirical evaluation plan (guidelines). Because IoTMindCare is a reference architecture, empirical claims about reductions in FPR/FNR require the instantiation of prototypes and controlled evaluation. For future prototypes we recommend the following evaluation regimen: (1) choose scenario(s) from Table 9.7; (2) run pre/post personalization experiments where the baseline uses population thresholds and the experimental arm uses IoTMindCare personalization and feedback loops; (3) report FPR, FNR, sensitivity, specificity, detection latency, escalation latency, and calibration (Brier score); (4) include robustness tests: sensor dropout (remove sensor modalities), concept-drift injection (temporal distribution shifts), and realistic network outages; (5) present results as absolute and relative changes in FPR/FNR or bootstrap estimates to quantify uncertainty.

9.5.3 Limitations and Threats to Validity

We summarize key limitations of *IoTMindCare* as follows, adapted from standard research design frameworks such as the SPHERE deployment framework [169]. These include threats

to validity in system evaluation, deployment feasibility, and ethical compliance. While our work lays a conceptual and architectural foundation, future empirical studies are required to address these threats.

- **Construct Validity:** Behavioral and physiological signals may not fully capture the complex psychological aspects of depression and have not been clinically validated against gold-standard diagnostic tools.
- **External Validity:** The system has not yet been deployed across diverse demographics, socioeconomic groups, or cultural contexts, limiting generalizability.
- **Ethical and Legal Constraints:** The system has not undergone formal ethical review or compliance assessment with regulations such as GDPR or HIPAA, especially regarding passive data collection.
- **Technical and Ecological Validity:** Assumes consistent sensor reliability, user compliance, and stable network conditions; real-world variability and behavioral shifts due to monitoring awareness may affect performance.

The validity issues include methodological concerns, such as the lack of clinical validation and the absence of empirical evaluation, as well as practical limitations in real-world deployment and ethical compliance. Specific challenges include construct and internal validity gaps, potential algorithmic bias, and uncertainties in scalability under resource constraints. Although the architecture presents a promising direction for intelligent depression care, future work must empirically validate its components, address ethical safeguards, and ensure robustness across user populations.

9.6 Conclusions

This research proposes a modular, layered IoT-based reference architecture (*IoTMindCare*) for continuous and personalized depression care in non-clinical settings. Integrating multi-modal sensing, edge-cloud analytics, adaptive feedback loops, and safety-driven design principles, our architecture addresses key challenges identified in existing literature, particularly safety, personalization, real-time response, and architectural extensibility. Through comparative analysis and design modelling, we demonstrated how *IoTMindCare* offers a comprehensive framework to guide the development of intelligent mental health monitoring systems.

While *IoTMindCare* offers a strong conceptual and technical foundation, several opportunities remain for future advancement. Real-world deployment and clinical trials would

be instrumental in validating the system's effectiveness in diverse populations. Adaptive learning components may be incorporated to support long-term personalization and reduce false positive or negative detections. Additionally, integration with healthcare infrastructures such as electronic health records and alignment with privacy and regulatory standards could support broader adoption. Moreover, the prototype can be strengthened by simulating a real-world environment. These directions represent promising avenues for translating IoT-based mental health architectures from prototype to practice.

Abbreviations

The following abbreviations are used in this manuscript:

Table 9.8 List of Abbreviations

Abbreviation	**Definition**
IoT	Internet of Things
ADD	Attribute-Driven Design
ATAM	Architecture Trade-off Analysis Method
MQTT	Message Queuing Telemetry Transport
CoAP	Constrained Application Protocol
UML	Unified Modeling Language
ADL	Activities of Daily Living
HAR	Human Activity Recognition
AAL	Ambient Assisted Living
QAS	Quality Attribute Scenario
FPR	False Positive Rate
FNR	False Negative Rate

Chapter 10

Prelude to the Expert-Informed Refinement of IoTMindCare

10.1 Introduction

Chapters 1–9 developed the foundation for *IoTMindCare* by examining the main components required for safety-aware and personalised IoT-based depression management. The systematic literature review in Chapter 3 identified fragmentation in existing IoT-based depression and emotional well-being systems, particularly regarding system integration, personalisation, safety, interoperability, and implementation considerations. Chapter 5 examined how environmental data may provide an additional interpretive layer for understanding mental health-related patterns. Chapter 7 then extended this direction by exploring how users may interact with the system through game-based tasks, adaptive self-reporting, and behavioural feedback. Chapter 9 brought these components together by presenting *IoTMindCare* as an integrative reference architecture for IoT-based depression management.

A central challenge for IoT-based depression management is how to organise sensing, analytics, calibration, personalisation, alert handling, clinical oversight, interoperability, and privacy within a coherent system architecture. This is particularly important because depression-related behavioural indicators are often individualised, sensor data may be incomplete or noisy, and automated alerts may create risks if they are poorly calibrated or insufficiently governed. Chapter 11 builds on the *IoTMindCare* architecture presented in Chapter 9 by refining it through expert feedback from both technical and clinical perspectives.

Chapter 11 is positioned as the expert-informed refinement stage of the thesis. Rather than focusing on a single component, this chapter considers the architecture as a whole and examines how sensing, analytics, calibration, edge-cloud processing, alert governance,

clinical oversight, interoperability, and privacy should be organised within a coherent system design. The purpose of Chapter 11 is to strengthen the architecture by examining how it can better support personalisation, safety, governance, and deployment-oriented design. This expert-informed stage should be understood as assessing the perceived suitability of architectural design decisions and informing their refinement, rather than establishing empirical evidence of system performance, clinical effectiveness, algorithmic validity, long-term user acceptance, or real-world deployment outcomes.

Chapter 11 also strengthens the thesis by incorporating multidisciplinary expert perspectives. Technical experts provide feedback on issues such as multimodal sensing, model initialisation, anomaly detection, calibration, edge-cloud allocation, and cloud integration. Clinical experts provide feedback on alert reliability, patient involvement, workflow integration, trust, privacy, consent, and safety. These perspectives are important because IoT-based depression management is not only a technical problem. It also involves clinical judgement, user burden, ethical responsibility, and governance requirements.

The expert evaluation involved 22 participants, comprising 14 technical and 8 clinical and healthcare-related experts. The technical panel represented backgrounds in IoT, software and AI engineering, cloud infrastructure, networking, biomedical engineering, research and development, and digital health. The clinical panel included expertise in psychiatry, general practice, mental health nursing, therapy research, clinical physiology, and healthcare safety. Participants were recruited purposively through professional and academic networks based on the relevance of their expertise to the proposed architecture.

Expert engagement was conducted through scenario-based questionnaires addressing the architecture's technical and clinical requirements. Participants evaluated proposed capabilities, including sensing, personalisation, edge-cloud processing, alerting, clinical workflows, and governance. The evaluation captured professional judgements regarding selected architectural functions and design decisions.

The expert evaluation reported in Chapter 11 is supported by thesis appendices that provide the associated ethics and participant-facing materials. The ethics approval documentation is provided in Appendix A, the participant information sheet is provided in Appendix B, the questionnaire is provided in Appendix C, and the feedback reports provided to participants following the questionnaire are included in Appendix D.

10.2 Contribution to the Thesis

The main contribution of Chapter 11 is the expert-informed refinement of *IoTMindCare* as a proposed safe and personalised IoT-based architecture for depression management.

Chapter 11 shows that an architecture for this context must go beyond collecting sensor data and applying machine learning models. It must also define how behavioural changes are calibrated, interpreted, filtered, escalated, reviewed, and governed. More specifically, Chapter 11 contributes to the thesis in three main ways:

- A dedicated analytics and calibration layer: This layer makes explicit the functions required for personalised baseline formation, hybrid model initialisation, multimodal fusion, anomaly detection, context-aware interpretation, threshold tuning, and ongoing recalibration. This refinement is important because the previous chapters showed that depression-related signals are complex, individualised, and context-dependent. Therefore, the architecture needs mechanisms that support interpretation and personalisation, rather than relying only on generic prediction outputs.
- The hybrid edge-cloud processing view: This view clarifies how responsibilities can be distributed between local edge resources and cloud infrastructure. Edge processing is positioned as important for low-latency screening, short-term interpretation, immediate alerting, and local resilience. Cloud processing is positioned as important for longitudinal storage, model retraining, drift detection, interoperability services, and governance support. This helps connect technical processing decisions with the broader requirements of safety, scalability, privacy, and clinical use.
- Traceable mapping between expert findings and architectural refinements: This mapping shows how expert feedback informed specific changes to the architecture. Therefore, it strengthens the transparency of the design process and demonstrates that the refined architecture is not a set of disconnected technical features, but a structured response to the requirements identified across the thesis.

10.3 Connection to Subsequent Chapters

Chapter 11 provides the final refinement of *IoTMindCare* before the thesis discussion and conclusion in Chapters 12 and 13. By incorporating technical and clinical expert perspectives, it clarifies how the architecture can better address personalisation, calibration, alert governance, clinical oversight, interoperability, privacy, and deployment-oriented design.

These refinements support the final discussion of the thesis contribution by showing how *IoTMindCare* responds to the limitations identified in earlier chapters. At the same time, the chapter also helps define the boundaries of the research. Although it strengthens the architecture through expert-informed refinement, empirical effectiveness, long-term user

acceptance, real-world deployment performance, and clinical validation remain areas for future implementation and evaluation.

Chapter 11

An Expert-Informed Architecture for Deployment-Oriented IoT-Based Depression Management

11.1 Abstract

The use of Internet of Things (IoT) technologies for depression management has increased interest in continuous behavioural monitoring using smartphones, wearables, and ambient sensors. However, many existing studies remain focused on predictive performance at the algorithmic level, with less attention to how personalisation, calibration, alert governance, hybrid edge-cloud processing, and deployment-related constraints should be integrated into a coherent system architecture. This study presents an expert-informed refinement of IoT-MindCare, an IoT-based architecture for depression management, with particular emphasis on personalisation, safety, and deployment-oriented design. A cross-sectional mixed-method expert evaluation was conducted with 22 domain experts, including 14 technical experts and 8 clinical experts. Quantitative responses were analysed using descriptive statistics and percentage agreement, while qualitative responses were analysed using inductive thematic analysis. The findings show strong expert support for multimodal sensing, edge-based analysis, cloud retraining, user feedback, and cloud logging within the technical panel, alongside strong clinical support for multimodal monitoring, personalised recommendations, and the importance of reducing alert delays and false positives. The results also highlight key design requirements including personalised baseline formation, hybrid model initialisation, multimodal fusion, context-aware interpretation, persistence windows, cooldown mechanisms, escalation logic, clinical oversight, and governance-oriented interoperability.

In response, the architecture was refined through the introduction of a dedicated analytics and calibration layer and a hybrid edge-cloud processing view. The study contributes a traceable, deployment-oriented architecture that moves beyond algorithm-centred design and provides clinically and technically informed guidance for IoT-based depression monitoring and decision support.

Publication Note

This chapter is based on the submitted article:

Zamani, S., Sinha, R., Madanian, S., & Nguyen, M. An expert-informed architecture for deployment-oriented IoT-based depression management. Submitted to Journal of Healthcare Informatics Research

The remainder of this chapter is based on the published article, with minor formatting adjustments made to align it with the structure and argument of this thesis.

11.2 Introduction

The rapid development of IoT technologies has expanded continuous health monitoring beyond conventional clinical settings. In digital mental health, smartphones, wearables, and ambient sensors are increasingly used to capture behavioural and physiological signals associated with depression [144, 139]. These data sources support passive sensing of sleep, activity, mobility, communication behaviour, and cardiovascular indicators such as heart rate and heart rate variability, enabling longitudinal assessment with minimal burden on users [204, 205].

Compared to periodic self-report and clinical observation, passive sensing can identify subtle behavioural changes that may precede deterioration. However, real-world deployment remains challenging. Behavioural anomaly detection systems may generate false-positive alerts when normal behavioural fluctuations are misclassified as clinically relevant changes [206]. Sensor unreliability, missing data, device limitations, and network interruptions can also degrade monitoring performance and system stability [207]. Clinically, such systems must also support trustworthy alerts, workflow integration, patient engagement, and regulatory compliance, including privacy, consent, and medical device requirements [208].

To address these challenges, this study presents an expert-informed architecture for IoT-based depression management, extending IoTMindCare [26] through input from both technical and clinical experts involved in mental health assessment and care delivery. The

refined architecture introduces a dedicated analytics and calibration layer to support personalised baseline formation, hybrid model initialisation, multimodal fusion, anomaly detection, and context-aware interpretation of behavioural change. In parallel, it defines a hybrid edge-cloud processing view that allocates low-latency screening and alerting functions to edge resources while assigning long-term storage, model adaptation, interoperability services, and governance support to cloud resources. Together, these refinements improve architectural coherence, deployment feasibility, and clinical relevance.

The main contributions of this study are as follows:

- Expert-informed technical, clinical, and operational requirements for IoT-based depression management through a structured expert evaluation.
- Refined logical architecture for IoTMindCare incorporating high-priority features identified by clinical and technical experts, such as a dedicated analytics and calibration layer supporting personalised baselines, context-aware interpretation, multimodal fusion, and calibration-driven decision-making [26]. The resulting logical architecture view shows the system's functional layering and cross-cutting concerns.
- Hybrid edge-cloud processing view to allocate responsibilities across edge and cloud resources to support timely screening, progressive personalisation, scalable retraining, and deployment-oriented governance requirements.
- Governance- and clinically-informed design guidance integrating safety, clinical oversight, privacy, interoperability, and regulatory considerations into the refined architecture to support future healthcare implementation.

The literature review in Section 11.3 identifies the gap in deployment-oriented architectural integration for IoT-based depression management, motivating the proposed expert-informed architecture. Section 11.4 presents the mixed-method expert evaluation design used to derive technical, clinical, and operational requirements. Section 11.5 presents the expert findings and introduces the refined logical architecture and hybrid edge-cloud processing view of IoTMindCare. Section 11.6 discusses the architectural implications of the findings, including calibration-driven modelling, alert governance, hybrid processing allocation, and governance-aware design. Finally, Section 11.7 summarises the main contributions of the study, while Section 11.7 outlines directions for future research and implementation.

11.3 Literature Review

This section reviews the literature relevant to the development of deployment-oriented IoT-based depression management systems. The review is organised progressively from behavioural sensing and personalised modelling to anomaly detection, alert governance, edge-cloud processing, and regulatory considerations. This structure reflects the study's architectural focus by moving from sensing and modelling foundations to system-level integration, deployment, and governance requirements. The section concludes by identifying the research gap that the proposed expert-informed IoTMindCare architecture addresses.

11.3.1 IoT-based Behavioural Sensing for Mental Health

IoT-enabled digital phenotyping has emerged as an effective approach for capturing behavioural markers associated with mental health conditions. Smartphones, wearables, and ambient sensors can continuously collect signals related to mobility, sleep, physical activity, communication behaviour, and physiological factors [209]. Recent reviews show that these modalities are among the most frequently studied indicators in digital psychiatry and may provide meaningful proxies for depressive symptom change [14, 210].

Empirical studies further indicate that depression is associated with measurable changes in everyday behaviour, including reduced mobility, irregular sleep patterns, and decreased activity levels [211, 212]. Nevertheless, the evidence base remains uneven. Recent systematic reviews report moderate predictive performance, substantial heterogeneity in feature engineering, and limited external validation, which restricts the clinical reliability of current models [13]. In practice, sensor noise, missingness, and device variability remain major obstacles to robust longitudinal monitoring.

11.3.2 Personalised Behavioural Modelling

A consistent finding in the literature is that behavioural patterns linked to depression are highly individualised [213]. Consequently, deviations from a person's own baseline are often more informative than comparisons with population-level thresholds. This has led to growing interest in idiographic and personalised modelling approaches.

Recent longitudinal evidence supports this direction. Balliu et al. [213] showed that personalised models built from smartphone behavioural data improved mood prediction substantially over non-personalised baselines. At the same time, other work has shown that passive sensing does not always outperform self-report in heterogeneous cohorts [214].

These findings indicate that reliable monitoring systems should combine passive sensing with personalised calibration rather than treating either component in isolation.

11.3.3 Machine Learning for Behavioural Anomaly Detection

Machine learning methods are widely used to identify behavioural deviations that may indicate deterioration in mental health. Because clinically labelled datasets are limited, many studies rely on unsupervised or semi-supervised learning, including clustering, autoencoders, one-class classifiers, and anomaly detection methods [215].

The literature consistently suggests that multimodal models are more informative than single-sensor approaches. Moshe et al. [211] found that combining smartphone, wearable, and self-reported data improved predictive performance relative to individual modalities. However, not all features contribute equally. Andersen et al. [212] reported that sleep-related variables were more informative than generic smartphone-use measures in predicting stress and depressive symptoms. Overall, the evidence supports modelling strategies that emphasise clinically meaningful features, temporal structure, and multimodal integration.

11.3.4 Alert Governance, Trust, and Workflow Integration

For clinically oriented monitoring systems, predictive performance is insufficient to ensure practical usefulness. Monitoring outputs must be translated into alerting and review processes that minimise unnecessary burden, preserve trust, and align with clinical workflows. In particular, the value of behavioural anomaly detection depends on whether changes can be identified, and also on how alerts are filtered, prioritised, reviewed, and escalated in practice [216].

This issue is especially important in mental health settings, where false positives may reduce user and clinician trust, while missed critical events may have serious consequences. Accordingly, deployment-oriented monitoring systems should incorporate persistence windows, alert suppression or cooldown mechanisms, contextual interpretation, and clinician review to help distinguish short-lived behavioural fluctuations from clinically meaningful change while limiting unnecessary alert burden [217]. These requirements should not be treated merely as operational considerations arising after model development; rather, they represent architectural design requirements that shape how sensing outputs are interpreted, how services coordinate alerting and review, and how responsibilities are distributed across users, caregivers, and clinicians. These considerations are consistent with broader digital mental health and healthcare governance literature emphasising safety, interpretability, human oversight, and responsible deployment [208, 218].

11.3.5 Edge-Cloud Architectures in IoT Healthcare

IoT healthcare systems increasingly adopt edge-cloud architectures to balance responsiveness, scalability, and privacy [219]. Edge devices enable local preprocessing and low-latency inference close to the data source, while cloud infrastructure supports long-term storage, model training, and cross-user analysis [164]. This division leads to continuous mental health monitoring, where timely detection and efficient handling of sensitive data are both required.

Recent reviews of IoT healthcare indicate that hybrid edge-cloud designs are well-suited to continuous monitoring applications, especially when data streams are high-volume, time-sensitive, and privacy-sensitive [219]. For behavioural monitoring systems, edge processing can reduce communication overhead and support early detection, while cloud-based analytics can facilitate personalised model adaptation and broader longitudinal learning [164]. Importantly, in such systems, edge-cloud design should not be viewed only as an infrastructure choice. Rather, it represents an architectural allocation of responsibilities across distributed resources, determining where preprocessing, feature extraction, anomaly screening, alert generation, storage, retraining, interoperability, and governance-related services should reside. This architectural perspective is particularly relevant to depression monitoring, where latency, resilience, privacy, and clinical accountability must be considered together.

11.3.6 Regulatory and Interoperability Considerations

Deployment in healthcare settings requires compliance with regulatory, interoperability, and governance requirements. Standards such as HL7 FHIR [220] support the structured exchange of health data and facilitate integration with electronic health records. In contrast, privacy frameworks such as GDPR [221] and HIPAA [222] impose important constraints on the collection, storage, and processing of sensitive health information. In addition, software used in clinical contexts may fall under medical device software requirements, particularly when monitoring outputs contribute to decision-making or intervention workflows.

Accordingly, systems intended for mental health monitoring must incorporate privacy, security, interoperability, and lifecycle management into the design stage rather than as post hoc additions [218]. These requirements are especially vital when monitoring outputs may influence clinical judgement, trigger alerts, or support ongoing care. While international frameworks provide useful reference points, real-world deployment also depends on jurisdiction-specific privacy, consent, data residency, and clinical governance requirements.

11.3.7 Research Gaps

Although the literature has established the promise of passive sensing and machine learning for mental health monitoring, most studies remain focused primarily on predictive performance at the algorithmic level. Comparatively less attention has been given to how personalised calibration, baseline formation, context-aware interpretation, alert governance, workflow integration, hybrid edge-cloud allocation, and deployment governance should be integrated into a coherent system architecture. As summarised in Table 11.1, prior work has advanced multimodal sensing and personalised modelling, but has given more limited attention to calibration-driven system design, architecture-level integration, and deployment-oriented workflow and governance requirements.

The resulting gap is therefore not simply the absence of more accurate predictive models. Rather, it is the limited treatment of architecture as the level at which behavioural sensing, personalisation, alert handling, clinical oversight, distributed processing, and governance constraints are brought together into an integrated design. Without such architectural integration, even promising sensing and modelling techniques remain difficult to translate into clinically credible and deployment-oriented systems.

This gap motivates the refinement of IoTMindCare as an expert-informed architecture for IoT-based depression management [26]. In response, this study extends the architecture through a dedicated analytics and calibration layer and a hybrid edge-cloud processing view that together support personalised baseline formation, multimodal interpretation, alert governance, and deployment-oriented clinical and governance requirements.

11.4 Research Design

To develop and evaluate an expert-informed architecture for deployment-oriented IoT-based depression management, this study employed a cross-sectional mixed-method expert evaluation design. Such approaches are well-suited to healthcare technology research when real-world deployment evidence is limited and informed expert judgement is needed to assess architectural relevance, perceived feasibility, clinical appropriateness, and deployment-oriented design requirements [224]. In this study, the evaluation collected both quantitative ratings and qualitative feedback to guide the evolution of the proposed architecture from both technical and clinical perspectives.

To develop and evaluate an expert-informed architecture for deployment-oriented IoT-based depression management, this study employed a cross-sectional mixed-method expert evaluation design. Such approaches are well-suited to healthcare technology research when real-world deployment evidence is limited and informed expert judgement is needed to

Table 11.1 Representative studies on IoT-based behavioural monitoring for mental health in relation to key architectural dimensions

Study	Multimodal sensing	Personalisation	Calibration	Architecture	Governance	Main limitation in relation to this study
[211]	Strong	Partial	Limited	Limited	Limited	Demonstrated benefits of multimodal fusion, but did not address calibration-driven design or deployment-oriented architecture.
[212]	Moderate	Limited	Limited	Limited	Limited	Identified informative behavioural features, especially sleep-related indicators, but lacked system-level integration and broader deployment considerations.
[214]	Moderate	Partial	Limited	Limited	Limited	Compared passive sensing with self-report, but did not translate these findings into an adaptive architectural framework.
[210]	Strong	Partial	Limited	Limited	Limited	Highlighted the promise of multimodal sensing, but also revealed heterogeneity and limited standardisation across studies.
[13]	Moderate	Limited	Partial	Limited	Partial	Identified missing data and validation challenges, but did not propose an architecture-level response oriented toward deployment.
[213]	Limited	Strong	Partial	Limited	Limited	Demonstrated the value of personalised modelling, but did not address cold-start handling, multimodal calibration, or hybrid deployment.
[223]	Partial	Strong	Partial	Partial	Partial	Combined generic and personalised modelling, but gave limited attention to calibration workflows, alert governance, and clinically oriented system integration.

Note: The ratings in this table are interpretive and indicate the relative extent to which each study addressed a given dimension. *Strong* indicates explicit and substantial treatment as a central component of the study; *Moderate* indicates meaningful but not comprehensive treatment; *Partial* indicates limited or supporting treatment; and *Limited* indicates minimal or no substantive treatment in relation to the architectural focus of the present study.

assess architectural relevance, perceived feasibility, clinical appropriateness, and deployment-oriented design requirements [224]. In this study, the evaluation collected both quantitative ratings and qualitative feedback to support the development of the proposed architecture. In particular, the study examined how IoTMindCare [26], initially developed from a predominantly technological perspective, could be extended through expert input to better address the practical, clinical, governance, and deployment-related requirements of depression management.

The study was conducted in the context of IoTMindCare [26], an IoT-based depression management architecture integrating multimodal sensing, hybrid edge-cloud processing, and anomaly detection capabilities. Given the safety-critical and clinically sensitive nature of depression management, the evaluation was designed around two complementary expert panels. The technical expert panel examined architectural components and design decisions related to sensing modalities, calibration, edge-cloud processing, anomaly detection, cloud integration, and regulatory considerations. The clinical expert panel examined design issues related to clinical usefulness, alert reliability, workflow integration, patient engagement, safety, and compliance with privacy and medical device requirements. Together, these panels provided complementary technical and clinical perspectives on the suitability of the architecture for deployment-oriented depression management.

A total of 22 experts participated in the study, comprising 14 technical experts and 8 clinical experts. Experts were recruited using purposive sampling based on their professional and academic experience relevant to the study aims. Potential participants were identified through professional networks, academic contacts, and direct invitations and invited to participate if they had substantial experience in fields aligned with the questionnaire domains. Eligibility for the technical panel required expertise in areas such as IoT systems, machine learning, software architecture, digital health technologies, or related engineering and computing domains. Eligibility for the clinical panel required relevant experience in mental health practice, psychological care, clinical decision-making, or related healthcare roles. Across both panels, participants were selected based on domain knowledge and their capacity to provide informed judgment on architectural feasibility, clinical relevance, safety, and deployment-oriented design requirements. Figure 11.1 illustrates the workflow of the expert evaluation process, from expert recruitment and questionnaire design to data collection, analysis, and architecture refinement.

Two domain-specific questionnaires were developed for technical and clinical experts. The instruments were aligned with each panel's expertise while retaining a comparable overall structure to support cross-group interpretation. The questionnaires were organised around key architectural and deployment-oriented themes, including data collection, system initialisation,

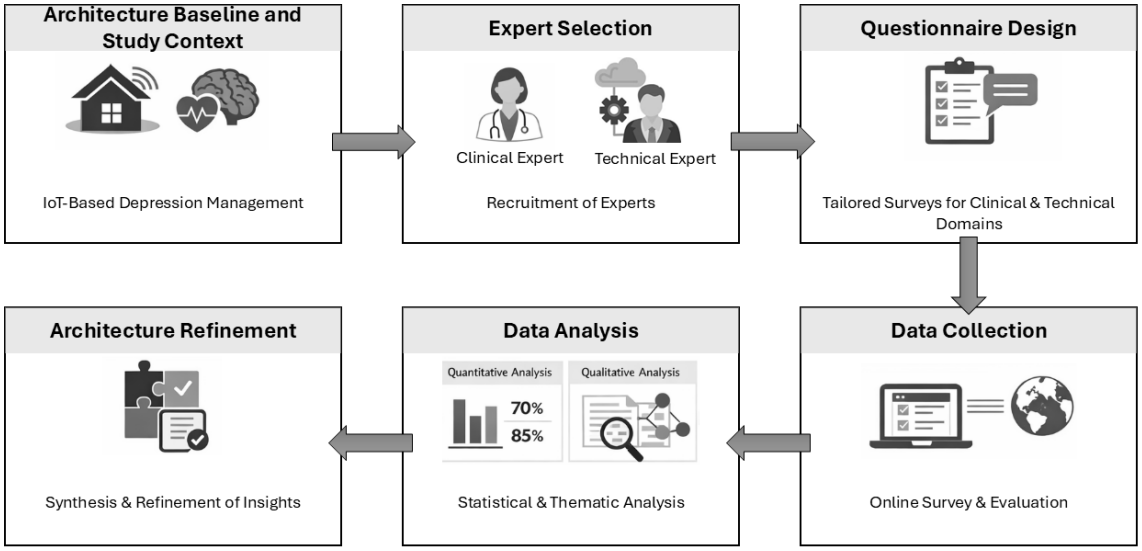


Fig. 11.1 Workflow of the mixed-method expert evaluation, showing the initial architecture context, expert recruitment, questionnaire design, data collection, analysis, and architecture refinement.

data processing, monitoring mechanisms, anomaly detection, feedback processes, technical or clinical practice context and perceptions, and legal and regulatory considerations. Each questionnaire included Likert-scale items to assess agreement and perceived importance, ranking and multiple-choice questions to prioritise system components, matrix-based items to map functionalities to architectural concerns, and open-ended questions to capture more detailed expert perspectives. The questionnaires were pilot-tested before distribution to improve clarity, completeness, and relevance. The technical panel included researchers and practitioners with experience in IoT, machine learning, software systems, and digital health, while the clinical panel included professionals with experience in mental health assessment, care delivery, and patient management.

Data were collected through a secure online survey platform, enabling geographically distributed participation while preserving anonymity. Experts evaluated key aspects of the proposed architecture, including sensing modalities, personalisation strategies, edge-cloud processing, anomaly detection mechanisms, cloud services, alert workflows, clinical role allocation, and regulatory considerations.

Quantitative responses were analysed using descriptive statistics, including mean, median, standard deviation, and percentage agreement. For Likert-scale items, agreement was defined as the proportion of respondents selecting *agree* or *strongly agree*. Neutral responses were retained in the denominator but were not counted as agreement. A 70% agreement threshold was used as an interpretive indicator of substantial expert support, consistent with prior

consensus-methods literature showing that percentage-based agreement thresholds in Delphi [225] and related expert panel studies commonly fall in the 70-80% range, with 75% reported as the median threshold in a systematic review of Delphi studies [226–228]. Values above 85% were interpreted in this study as indicating stronger support; however, this higher threshold was used only as an interpretive guide rather than as a universal or formal Delphi criterion. Accordingly, these thresholds were used to distinguish areas of relatively stronger and weaker expert endorsement, rather than to claim consensus in the strict Delphi sense, for which definitions remain heterogeneous across studies and should be specified a priori [226, 229, 230].

The qualitative component was analysed using reflexive thematic analysis to systematically identify patterns in expert feedback and inform architectural refinement. Qualitative responses from open-ended questions were analysed using reflexive thematic analysis following Braun and Clarke's six-phase framework [231]. The analysis began with data familiarisation, followed by inductive coding to identify meaningful patterns across expert responses without imposing predefined categories. Codes were then grouped into candidate themes representing key design requirements, concerns, and priorities.

Themes were iteratively reviewed and refined to ensure coherence and alignment with the original data, and subsequently defined in relation to architectural, clinical, and deployment considerations. An iterative and reflexive approach was adopted throughout the process to enhance analytical trustworthiness. Theme frequency was reported as the number of participants referencing a given idea; as responses could contribute to multiple themes, counts are not mutually exclusive.

Findings from the quantitative and qualitative analyses were integrated to identify key areas of support, highlight critical design requirements, and reveal important limitations or concerns. Rather than treating the exercise as a formal consensus-building process, the analysis was used to support evidence-informed architectural refinement, while also retaining less common but important concerns where these had implications for safety, clinical use, or deployment. To support traceable refinement, the key findings were mapped to specific architectural implications and then translated into changes affecting either the logical architecture view, the hybrid edge-cloud processing view, or both.

Ethical approval was obtained from the Auckland University of Technology Ethics Committee (AUTEC). Participation was voluntary, informed consent was obtained from all experts, and all responses were anonymised and reported in aggregated form.

Overall, this methodology provides a structured multidisciplinary basis for expert-informed architectural refinement of IoTMindCare [26], supporting assessment of its technical relevance, clinical appropriateness, and deployment-oriented design considerations.

11.5 Results

This section presents findings from technical and clinical expert evaluations and explains how they informed refinement of the IoTMindCare architecture [26]. The results are organised by expert panel and, within each panel, into qualitative and quantitative components. The qualitative analyses identify the principal design themes related to calibration, model initialisation, edge-cloud allocation, anomaly detection, clinical workflow, alert trust, and regulatory requirements, while the quantitative results summarise the level of expert agreement supporting these design directions.

11.5.1 Findings from Technical Expert Feedback

Qualitative Analysis

Calibration and system initialisation The qualitative analysis of open-ended responses from technical experts showed strong support for multi-layer calibration across sensing, modelling, and alert management. As summarised in Table 11.2, the most frequently cited requirements were personalised behavioural baselines, sensor calibration and quality assurance, threshold tuning with temporal persistence, context-aware interpretation, and multi-sensor fusion.

This feedback suggests that reducing false positives depends on the interaction among individual behavioural patterns, contextual information, and robust sensing, rather than on threshold tuning alone. Experts also emphasised the handling of missing data and the control of alert burden, indicating that system reliability requires mechanisms capable of distinguishing transient behavioural variation from clinically meaningful change. Architecturally, these findings support the introduction of a dedicated analytics and calibration layer responsible for personalised baseline formation, threshold tuning, multimodal fusion, context-aware interpretation, and calibration-driven alert control.

Model initialisation Experts consistently supported a hybrid model initialisation strategy. Personalised models were considered essential for capturing individual behavioural variability and improving detection accuracy. In contrast, generic pre-trained models were viewed as necessary during the onboarding phase to address the cold-start problem.

Overall, the findings in Table 11.2 indicate that population-level models should provide an initial stable baseline, followed by progressive personalisation as user-specific data accumulate. This staged approach balances early-stage robustness with longer-term sensitivity to within-person behavioural change. This supports explicitly representing hybrid initialisation

Table 11.2 Technical Expert Feedback: Calibration, personalised initialisation, and generic model setup

Key parameters for system calibration to reduce false positives		
Theme	**Description**	**Frequency**
Personalised baselines	Calibration of individual behavioural norms rather than population averages.	4
Sensor calibration and quality assurance	Device placement, sampling configuration, and noise filtering.	4
Thresholds, smoothing, and persistence	Temporal smoothing and persistence rules to avoid transient alerts.	4
Context-aware rules	Incorporation of behavioural and environmental context.	4
Multi-sensor fusion	Weighted integration of multiple sensing modalities.	4
Alert control mechanisms	Cooldowns, aggregation windows, and prompt limits.	3
Missing-data handling	Tolerance thresholds and imputation strategies.	2
Model-level tuning	Adjustment of anomaly thresholds and feature weights.	2
Clinical linkage	Mapping anomalies to clinical interpretation.	1
Feedback on Personalised Model Initialisation		
Theme	**Description**	**Frequency**
Individual variability	Behavioural indicators vary across users; personalisation improves accuracy.	5
Cold-start limitation	Lack of initial data limits early-stage performance.	3
Hybrid modelling	Combination of generic and personalised models recommended.	2
Overfitting risk	Sensitivity to short-term behavioural variation.	2
Operational complexity	Increased computational and maintenance requirements.	2
Clinical oversight	Need for clinician involvement in interpretation.	1
Expert Perspectives on Generic Pre-Trained Models		
Theme	**Description**	**Frequency**
Cold-start practicality	Enables system operation without user history.	2
Stable baseline	Provides reliable early-stage monitoring.	1
Noise robustness	Reduces sensitivity to limited or noisy data.	1
Progressive personalisation	Transition toward personalised models over time.	2

Note: Frequencies indicate the number of respondents who mentioned a theme. Because individual responses could refer to more than one theme, totals may exceed the number of participants.

and progressive adaptation within the logical architecture, while also informing the hybrid processing view, in which ongoing model updating and retraining are allocated across edge and cloud resources.

Edge-cloud processing Experts showed strong agreement that edge and cloud computing play complementary roles (Table 11.3). Edge processing was associated with low-latency detection, resilience, and privacy, whereas cloud platforms were linked to scalable computation, model retraining, and long-term analytics.

Table 11.3 Deployment architecture, cloud platform selection, and cloud integration

Edge vs cloud processing trade-offs		
Theme	**Description**	**Frequency**
Edge advantages	Low latency, offline capability, privacy preservation.	6
Edge limitations	Limited computation and storage.	5
Cloud strengths	Scalable computation, storage, retraining, analytics.	6
Hybrid recommendation	Edge for real-time, cloud for long-term processing.	5
Operational trade-offs	Cost, security, and connectivity considerations.	6
Cloud platform selection for IoT-based depression management		
Theme	**Short description**	**Frequency**
AWS / Amazon Web Services	Mature IoT ecosystem, edge support, and microservices for deployment.	9
Microsoft Azure / Azure IoT Hub	Strong healthcare integration, FHIR services, EHR support, and security.	6
Google Cloud / Vertex AI / GCP	Strong AI and analytics capabilities, but weaker native IoT support.	4
IoT-specific platform features	Sensor ingestion, dashboards, IoT hubs/core services, and edge computing support.	4
Machine learning / AI support	Cloud AI/ML tools for prediction and optimisation.	2
Real-time dashboards / ingestion	Practical support for live monitoring and data pipelines.	2
Scalability / deployment feasibility	Ability to support larger IoT deployments and production workloads.	2
Healthcare / clinical support	FHIR, EHR integration, and healthcare security requirements.	1
Cloud service integration and deployment constraints		
Theme	**Short description**	**Frequency**
Integrated multi-service architecture	Experts preferred a fully integrated and interdependent cloud stack.	2
Privacy and data residency	Clinical data should comply with local privacy and residency requirements.	1
Specific cloud services	Azure API Management, AWS API Gateway, Azure Health Data Services (FHIR).	1
Survey design limitation	Radio buttons restricted multiple selections.	1

Note: Frequencies indicate the number of respondents who mentioned a theme. Because individual responses could refer to more than one theme, totals may exceed the number of participants.

These findings support a hybrid processing allocation in which low-latency screening and local resilience are handled at the edge, while longitudinal analytics, retraining, and governance-related services are handled in the cloud. Rather than treating edge and cloud as interchangeable infrastructure options, the findings suggest that they should be understood as serving distinct but complementary architectural responsibilities.

Cloud platform selection and service integration Responses on cloud platform selection in Table 11.3 indicate a clear preference for hyperscaler ecosystems over lightweight proto-

typing platforms. AWS was the most frequently cited option, primarily because of its mature IoT ecosystem, flexible microservices support, and suitability for edge-enabled deployments. Microsoft Azure was also strongly supported, particularly for healthcare-oriented capabilities such as FHIR services, EHR integration, and security support. Google Cloud was mentioned less frequently and was mainly associated with AI and analytics capabilities, although several experts noted that its native IoT support is less comprehensive than that of AWS or Azure.

Responses emphasised that IoT-specific platform features, including sensor ingestion, dashboards, and edge computing support, are important for practical deployment. Overall, these findings suggest that cloud platform selection should be guided primarily by interoperability, integration capability, healthcare-oriented services, and deployment-oriented requirements. Within this expert sample, AWS and Azure were the most frequently preferred platforms, while Google Cloud was more often cited for specialised analytics and AI functions.

Responses further indicated that cloud services should be treated as part of an integrated and interdependent architecture rather than as isolated components. Experts emphasised the importance of privacy, data residency, and regulatory compliance for clinical data, and referred to services such as Azure API Management, AWS API Gateway, and Azure Health Data Services (FHIR) as useful examples for supporting interoperability and secure integration.

Taken together, these responses reinforce the need for a deployment-oriented cloud layer in IoTMindCare that supports secure API management, healthcare interoperability, and regulatory compliance. They also suggest that platform selection should prioritise integration readiness and clinical deployment considerations, with lightweight environments remaining more suitable for prototyping than for production-oriented use.

Anomaly detection Experts identified unsupervised and semi-supervised methods as the most appropriate approaches for behavioural anomaly detection, particularly in the absence of labelled clinical data (Table 11.4). Temporal models such as LSTM and GRU were recommended for capturing sequential behavioural patterns, while hybrid and ensemble approaches were preferred to improve robustness. Architecturally, this reinforces the need to treat anomaly detection as part of a broader calibration and interpretation pipeline rather than as an isolated model-selection problem.

The findings further indicate that feature selection, temporal persistence, and multimodal fusion are more critical than algorithm selection alone, reinforcing the need for calibration-driven modelling.

Table 11.4 Anomaly detection methods and regulatory/standards considerations

Machine learning approaches for anomaly detection		
Theme	**Description**	**Frequency**
Unsupervised methods	Isolation Forest, autoencoders, clustering.	5
Temporal models	LSTM/GRU for sequential behavioural patterns.	5
Supervised methods	Used when labelled data are available.	4
Hybrid approaches	Multi-model and multimodal fusion strategies.	3
Regulatory and standards considerations		
Theme	**Description**	**Examples**
Privacy and data protection	Consent, minimisation, secure processing	GDPR, HIPAA
Interoperability	Standardised data exchange	HL7 FHIR, SNOMED CT
Medical device compliance	Software lifecycle and safety standards	IEC 62304, ISO 13485
Security controls	Access control, encryption, auditing	ISO 27001
Clinical governance	Oversight and accountability mechanisms	Clinical protocols

Note: Frequencies indicate the number of respondents who mentioned a theme. Because individual responses could refer to more than one theme, totals may exceed the number of participants.

Regulatory considerations Regulatory and governance requirements were identified as integral to system design (Table 11.4). Experts emphasised privacy and data protection, interoperability standards, and medical device compliance as core requirements. Security controls, data residency, and clinical governance were also highlighted.

These findings indicate that compliance considerations must be embedded throughout the system lifecycle, from data collection and processing to alert generation and clinical integration. Architecturally, this supports representing governance and interoperability as cross-cutting concerns that constrain data handling, service integration, alert workflows, and clinical use across the system.

Quantitative Analysis

The closed-ended responses showed strong technical support for the proposed architecture across its core functions. Full agreement was observed for edge-based analysis, cloud retraining, user feedback, and cloud logging (14/14, 100.0%). High agreement was also reported for multimodal sensing, post-processing, and bidirectional communication (13/14, 92.9%), while edge feasibility, latency reduction, and overall technical feasibility each reached 12/14 (85.7%). Missing-data handling received the lowest level of agreement

among the closed-ended items (11/14, 78.6%), indicating that although experts recognised its importance, implementation choices in this area require careful design and validation. Overall, these findings provide substantial support for multimodal sensing, real-time edge processing, cloud-side retraining, and feedback-driven monitoring, while also indicating that missing-data handling should remain adaptive and context sensitive.

11.5.2 Findings from Clinical Expert Feedback

Qualitative Analysis

Clinical experts reported limited awareness of specific digital depression monitoring tools. Although a small number mentioned consumer applications such as Daylio [232], most respondents could not recall a specific system. This suggests that digital mental health tools are not yet strongly embedded in routine clinical awareness or practice.

Across the regulatory and ethical questions, privacy and informed consent emerged as the most consistent concerns. Experts highlighted the Privacy Act 2020 [233], medical device regulations, patient safety, cybersecurity, accessibility, and cultural considerations as central requirements. They also emphasised clinician accountability, escalation protocols, and ethical AI use, indicating that clinical deployment must preserve oversight, traceability, and equitable access. When the system was framed as a patient-facing support tool, experts supported a lighter regulatory burden, but still required strong safeguards for privacy, safety, and validated self-monitoring (Table 11.5).

Open-ended responses reinforced these perspectives. Experts described the proposed architecture as promising for early alerts, personalised support, and real-time monitoring, while also stressing the need to balance engagement, privacy, safety, and clinical accuracy. Several respondents noted that patients with severe depression may experience difficulties with adherence, making usability and low-burden design especially important. Personalisation was repeatedly identified as essential for sustained engagement and therapeutic alliance. Architecturally, these findings support a stronger emphasis on alert governance, human review, low-burden interaction, and role-aware workflow design within the service and interface layers.

Quantitative Analysis

The closed-ended responses showed strong clinical support for the proposed architecture. Full agreement was reached for multimodal data for anomaly detection, personalised recommendations, patient involvement in setup, digital literacy, and the overall value of digital health (8/8, 100.0%). Experts also unanimously agreed that delayed alerts and false positives

Table 11.5 Main themes from open-ended responses by clinical experts

Theme	**Key insight**	**Frequency**
Personalisation	Tailor recommendations and thresholds to the individual.	2
Trustworthy alerts	Reduce delay and false positives; keep alerts clinically reliable.	3
Role-specific workflows	Distribute responsibilities across clinical roles.	3
Regulatory compliance	Address privacy, consent, safety, and medical device rules.	6
Engagement and adherence	Minimise burden to support long-term use.	4

Note: Frequencies indicate the number of respondents who mentioned a theme. Because individual responses could refer to more than one theme, totals may exceed the number of participants.

erode trust (8/8, 100.0%), underscoring the importance of timely, clinically reliable alerting. Moderate agreement was observed for local alerts for mild or moderate cases, trust in alerts at $\geq$90% accuracy, event summaries via an app or portal, and the system's potential to address unmet needs (6/8, 75.0%), while patient confirmation or cancellation of false alarms received 7/8 (87.5%). In contrast, only 2/8 (25.0%) agreed that missing an alert may be acceptable, indicating strong expert rejection of false negatives in safety-critical situations. Overall, these findings suggest that clinical acceptance depends not only on personalisation and multimodal monitoring but also on high-sensitivity alerting, trustworthiness, and the preservation of human oversight.

11.5.3 Extending the IoTMindCare Architecture

The technical and clinical findings together informed two complementary refinements of the IoTMindCare architecture. First, the logical architecture was extended by introducing a dedicated analytics and calibration layer to support hybrid model initialisation, personalised baselines, multimodal fusion, anomaly detection, context-aware interpretation, and calibration-driven decision-making. In parallel, the service layer was refined to more explicitly reflect alert governance requirements, including persistence windows, cooldown mechanisms, escalation logic, clinician review, and workflow support. These refinements are reflected in the logical architecture view shown in Figure 11.2, which presents the system's functional layers alongside the cross-cutting roles of progressive personalisation, safety and clinical oversight, and governance and interoperability.

Second, the expert findings reinforced the need for a hybrid edge-cloud processing view in which computational responsibilities are distributed based on latency, scalability,

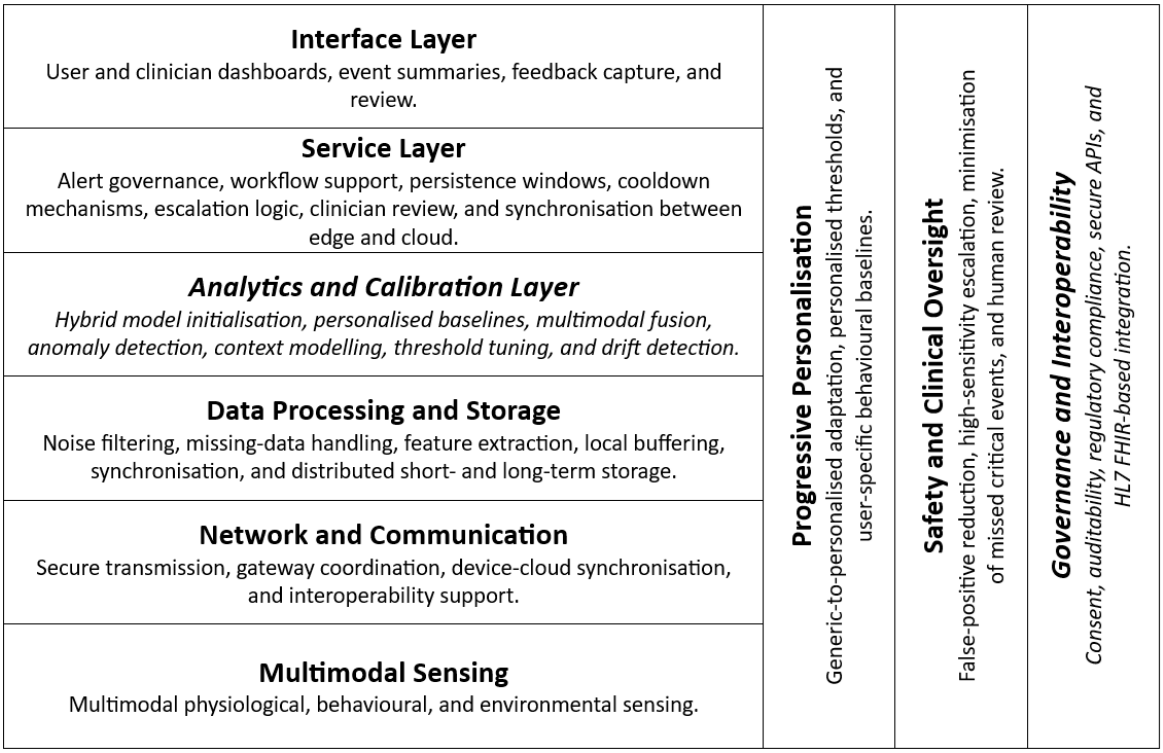


Fig. 11.2 Logical architecture view of IoTMindCare, showing the functional layering of the system and the cross-cutting roles of progressive personalisation, safety and clinical oversight, and governance and interoperability.

resilience, and governance requirements. In this refined view, edge resources support local preprocessing, feature extraction, short-term context interpretation, immediate screening, low-latency alerting, and offline resilience, while cloud resources support longitudinal storage, baseline updating, model retraining, drift detection, cross-user learning, interoperability services, and governance support, as shown in Figure 11.3.

Across both views, the findings also highlighted several cross-cutting design priorities. Progressive personalisation must move from generic model initialisation toward user-specific adaptation as more behavioural data becomes available. Safety and clinical oversight must address not only false-positive reduction but also the minimisation of missed critical events through high-sensitivity escalation and human review. In addition, governance and interoperability requirements, including privacy, consent, auditability, interoperability, and regulatory compliance, must be embedded throughout the architecture rather than treated as external constraints.

To make the refinement process explicit, Table 11.6 summarises how the main technical and clinical expert findings were translated into specific architectural refinements. This

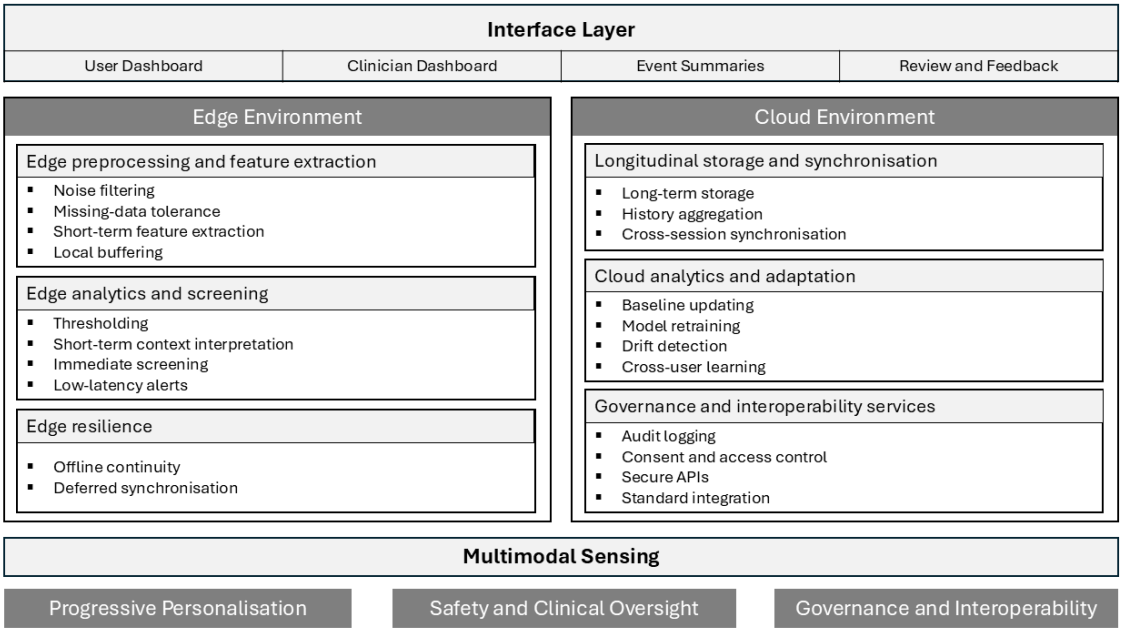


Fig. 11.3 Hybrid edge-cloud processing view of IoTMindCare, showing the allocation of real-time screening, local resilience, and immediate alerting to edge resources, and the allocation of longitudinal analytics, model adaptation, interoperability services, and governance support to cloud resources.

mapping provides traceability between the expert-derived requirements, the resulting design decisions, and the architecture views affected by those decisions.

11.6 Discussion

This study aimed to extend IoTMindCare [26], an IoT-based architecture for depression management through structured input from both technical and clinical experts. Overall, the findings suggest that a deployment-oriented depression monitoring architecture should extend beyond multimodal sensing and anomaly detection alone. Instead, it should explicitly incorporate personalised baseline formation, hybrid model initialisation, context-aware interpretation, alert governance, hybrid edge-cloud processing, and cross-cutting governance requirements. In this regard, the present work contributes an architecture-level response to a gap in the literature, where prior studies have focused primarily on predictive modelling performance rather than how to integrate them into a clinically and operationally coherent system.

One of the most important findings was the strong emphasis on calibration. Technical experts consistently highlighted the need for personalised baselines, sensor quality assurance,

Table 11.6 Traceable mapping from expert findings to architectural refinements

Expert finding	**Evidence source**	**Architectural refinement**	**Affected view**
Need for personalised baselines and threshold tuning	Section 11.5.1	Introduction of a dedicated analytics and calibration layer supporting baseline formation, threshold tuning, and calibration-driven decision-making	Logical view
Need for generic and personalised model initialisation	Section 11.5.1	Hybrid model initialisation with progressive personalisation as user-specific data accumulate	Logical and processing views
Need for multimodal fusion and context-aware interpretation	Section 11.5.1	Analytics and calibration layer expanded to include multimodal fusion and context modelling	Logical view
Need for low-latency screening and offline resilience	Sections 11.5.1, 11.5.1	Allocation of preprocessing, short-term interpretation, immediate screening, and local alerting to edge resources	Processing view
Need for retraining, longitudinal analytics, and drift detection	Sections 11.5.1, 11.5.1	Allocation of storage, retraining, drift detection, and cross-user learning to cloud resources	Processing view
Need to reduce false positives and alert burden	Sections 11.5.1, 11.5.2	Refinement of the service layer to include persistence windows, cooldown mechanisms, and alert governance logic	Logical view
Need to minimise missed critical events and preserve human oversight	Sections 11.5.2, 11.5.2	Safety and clinical oversight represented as a cross-cutting concern with escalation logic and clinician review	Both views
Need for privacy, consent, interoperability, and compliance	Sections 11.5.1, 11.5.2	Governance and interoperability represented as a cross-cutting concern affecting communication, storage, services, and integration	Both views

threshold tuning, temporal persistence, context-aware rules, and multimodal fusion. This indicates that behavioural anomaly detection in depression management should not be treated as a single model-selection problem. Rather, it should be understood as part of a broader calibration and interpretation pipeline in which sensing quality, behavioural context, temporal smoothing, and alert-control logic jointly affect reliability. This finding extends prior work on personalised behavioural modelling and passive sensing [213, 211, 214] by showing that the architectural mechanisms supporting personalisation are as important as the predictive models themselves.

A second key finding concerns the role of hybrid model initialisation. Experts supported the use of generic pre-trained models during onboarding together with progressive adaptation toward user-specific models as more data become available. This is especially important in mental health monitoring, where individual behavioural baselines differ substantially and where the cold-start problem can undermine early-stage performance. The implication is that personalisation should be designed as a staged architectural process rather than assumed to emerge automatically from downstream analytics. The introduction of a dedicated analytics and calibration layer therefore represents a meaningful refinement, because it makes explicit the functions required to support baseline formation, hybrid initialisation, anomaly detection, multimodal fusion, and ongoing recalibration.

The results also reinforce the value of a hybrid edge-cloud processing design. Technical experts associated edge resources with low-latency screening, local resilience, and privacy-preserving processing, whereas cloud resources were linked to longitudinal storage, model retraining, drift detection, interoperability, and governance-oriented services. This is consistent with broader IoT healthcare literature supporting edge-cloud complementarities in time-sensitive and privacy-sensitive environments [219, 164]. However, the present study extends that literature by mapping these processing responsibilities directly to the needs of depression monitoring, including alert timing, offline resilience, progressive personalisation, and governance support. Rather than treating edge–cloud design as generic infrastructure, the refined architecture links processing allocation to clinical and behavioural characteristics.

Clinical findings further highlighted that practical usefulness depends not only on detecting behavioural change, but also on how alerts are governed. Experts expressed strong concern about delayed alerts, false positives, and missed critical events, while also emphasising the need for clinician accountability, role-aware workflows, and patient involvement where appropriate. These findings suggest that alert governance should be treated as a first-class architectural concern. In this study, this led to refinement of the service layer to include persistence windows, cooldown logic, escalation pathways, clinician review, and workflow-aware alert handling. This strengthens the architecture by aligning technical monitoring outputs with the realities of clinical trust, human oversight, and operational burden. In other words, the architecture is not only detecting change, but also structuring how potentially meaningful change should be interpreted, prioritised, and acted upon.

Another important implication of the study is that governance and interoperability requirements should be embedded throughout the architecture rather than added as external constraints. Across both expert panels, privacy, consent, security, regulatory compliance, and interoperable data exchange were treated as integral to design. This aligns with growing calls in digital mental health for responsible, explainable, and governance-aware systems

[208, 218]. The refined architecture therefore represents these concerns as cross-cutting requirements affecting sensing, communication, analytics, services, and clinical integration. This is critical in mental health contexts, where sensitive data and clinical decisions demand traceability and accountability.

From a design perspective, the study makes three main contributions. First, it translates expert findings into a refined logical architecture in which analytics and calibration are made explicit rather than remaining implicit within service or data layers. Second, it proposes a hybrid edge-cloud processing view that allocates system responsibilities according to latency, resilience, scalability, and governance requirements. Third, it provides a traceable mapping from expert-derived requirements to architectural refinements. This traceability is important because it strengthens the transparency of architectural decision-making and improves the paper's value as a reusable design reference for related mental health monitoring systems.

11.6.1 Threats to Validity

This study has several threats to validity that should be considered when interpreting the findings.

Internal validity: The study relied on expert judgement to evaluate and inform the proposed architecture. While this approach is appropriate for early-stage architectural development, it introduces the possibility of subjective bias in responses and interpretation. In addition, the evaluation was conducted as a single-round expert study rather than a multi-round Delphi process, meaning that no iterative feedback or convergence of expert opinion was performed.

External validity: The expert sample was purposively selected and modest in size, comprising technical and clinical experts with relevant domain knowledge. As a result, the findings reflect informed professional judgement rather than statistical generalisability to a broader population of stakeholders or healthcare settings.

Construct validity: The study evaluated architectural requirements, perceived feasibility, and design appropriateness rather than real-world system performance. Therefore, the findings should be interpreted as evidence supporting architectural design decisions rather than validation of clinical effectiveness or predictive performance.

Contextual and ecological validity: Some regulatory and governance insights may be influenced by the professional and geographic context of the recruited experts. Consequently, real-world deployment would require adaptation to local legal, organisational, and clinical requirements, including jurisdiction-specific privacy, consent, and medical device regulations.

11.7 Conclusions and Future Work

This study presented an expert-informed refinement of IoTMindCare, an IoT-based architecture for depression management, with particular attention to personalisation, safety, and deployment-oriented design. Using a mixed-method expert evaluation involving both technical and clinical panels, the study identified key requirements for personalised baseline formation, hybrid model initialisation, multimodal fusion, context-aware interpretation, alert governance, hybrid edge-cloud processing, and governance-aware interoperability. These findings informed two complementary architectural refinements: the introduction of a dedicated analytics and calibration layer within the logical architecture, and the definition of a hybrid edge-cloud processing view that allocates responsibilities according to latency, resilience, scalability, and governance requirements.

The results show that a clinically relevant IoT architecture for depression management must go beyond sensor integration and predictive modelling alone. It must also address how behavioural deviations are calibrated, interpreted, reviewed, escalated, and governed in practice. In this sense, the contribution of the study lies not in claiming empirical effectiveness, but in providing a structured and traceable architecture refinement grounded in multidisciplinary expert judgement. The refined architecture offers a clearer foundation for future implementation and evaluation by explicitly integrating technical feasibility, clinical oversight, and governance considerations into the system design.

Overall, the study contributes an architecture-level perspective to the digital mental health literature and provides a deployment-oriented framework that can support future work on monitoring, decision support, and responsible IoT-enabled mental health systems.

Several directions emerge from this work as follows:

- Implement the refined architecture as a prototype to evaluate how the proposed logical and processing views operate under realistic sensing, communication, and alerting conditions, including latency, resilience, missing-data handling, interoperability, and workflow integration.

- Evaluate the architecture using longitudinal real-world data to test personalised baseline formation, hybrid model initialisation, drift detection, and alert-governance mechanisms over time, and to assess whether the calibration-driven design can reduce false positives while maintaining high sensitivity to clinically meaningful deterioration.

- Extend evaluation beyond expert judgement to include end users, caregivers, and service providers in order to better understand usability, burden, acceptability, trust, and

the practical implications of patient-facing event summaries, confirmation workflows, and clinician escalation pathways.

- Examine implementation across different regulatory and organisational settings, with attention to jurisdiction-specific privacy, consent, medical device, interoperability, and data residency requirements, in order to identify which architectural elements are transferable and which require local adaptation.
- Explore formal validation of traceability between stakeholder requirements and architectural decisions through methods such as requirement modelling, scenario-based testing, pilot deployment in clinical or community settings, and comparison with alternative digital mental health architectures.

Declarations

- Funding: No funding was received for conducting this study or for the preparation of this manuscript.

- Conflict of interest/Competing interests: The authors have no competing interests to declare that are relevant to the content of this article.
- Ethics approval and consent to participate: This study involved expert participants. Ethics approval was obtained from the Auckland University of Technology Ethics Committee (AUTEC) (Ethics approval number: 25/401; approval date: 09/12/2025). The questionnaire and methodology for this study were approved prior to data collection. All participants were provided with study information and gave informed consent before participating.

Chapter 12

Discussion

12.1 Introduction

This thesis addresses a central limitation in the current development of IoT-based systems for depression monitoring and management: the field has advanced technically but remains architecturally fragmented. Existing work shows that IoT devices, mobile applications, environmental sensors, self-report tools, and machine learning models can each contribute to emotional well-being and depression-related monitoring. However, the more difficult question is not whether these technologies can collect data or generate predictions. The more important question is how such technologies can be integrated into a safe, personalised, context-aware, and clinically meaningful system.

The central argument is that IoT-based depression management cannot be achieved through isolated sensing, generic prediction, or standalone application functionality. It requires an integrated architecture that connects continuous sensing, contextual interpretation, behavioural engagement, personalised analytics, governed escalation, human oversight, privacy, interoperability, and implementation-oriented design.

The chapter first addresses the research questions and then synthesises the thesis contribution through six thesis-level findings:

1. **IoT-based depression systems need to move from monitoring to management.** Detection alone is insufficient because depression support requires not only the identification of possible changes, but also interpretation, feedback, user engagement, calibration, and escalation. This matters because a technically accurate signal has limited value if it is not connected to a meaningful pathway for support or action.

2. **IoT-based depression systems need to move from generic prediction to personalisation.** Depression-related patterns vary between individuals, so fixed thresholds and

generic models may misinterpret normal variation as risk, or risk as normal variation. This matters because safe and meaningful support depends on personalised baselines, adaptive thresholds, and interpretation in relation to the user's own context and history.

3. **IoT-based depression systems need multimodal and context-aware interpretation.** No single data source can fully represent depression-related experience or behaviour. Environmental data, behavioural indicators, self-reports, and sensor-based data each provide partial evidence. This matters because the meaning of a signal often depends on its relationship with other signals and the context in which it occurs.

4. **IoT-based depression research needs to move from prototype functionality to architecture-level integration.** Individual functions such as sensing, prediction, dashboards, or alerts are useful, but they do not by themselves create a safe and coherent management system. This matters because future systems need an architecture that defines how components, data flows, stakeholder roles, safety mechanisms, personalisation, interoperability, and governance fit together.

5. **IoT-based depression systems need governed human-in-the-loop support rather than fully automated decision-making.** Automated systems may assist with monitoring and early warning, but depression management involves ambiguity, risk, personal context, and clinical judgement. This matters because safety depends on alert governance, human oversight, user control, caregiver boundaries, privacy, and appropriate clinical involvement.

6. **IoT-based depression systems need to move from technical feasibility to implementation-oriented design.** Demonstrating that data can be collected or analysed is not enough to support future use. This matters because systems intended for later piloting or evaluation need to consider edge-cloud processing, interoperability, governance, privacy, regulatory awareness, workflow integration, user acceptance, and clinical involvement from the design stage.

These findings provide the basis for the later synthesis section, following the discussion of the research questions. Each is explored through its connection to the corresponding research question, supported by findings from the thesis chapters, and considered in terms of its implications for IoT-based depression management as well as directions for further development and application.

12.2 Addressing the Research Questions

This section discusses how the thesis addresses the overarching research question and the four integrated research questions introduced in Chapter 1. The purpose of this section is to show how the chapters collectively move the field from fragmented IoT-based depression monitoring towards an integrated reference architecture for safe, personalised, context-aware, and future-adaptable system design.

12.2.1 Overarching Research Question

The overarching research question, introduced in Section 11.2, asked: *How can architecture-level requirements and design principles be derived and organised to address the fragmentation of current IoT-enabled depression monitoring and management approaches across sensing, interpretation, safety, personalisation, contextual awareness, engagement, stakeholder interaction, interoperability, and governance?*

The main finding made in response to this question is that these concerns cannot be treated as independent design issues. In the current state of the field, IoT-based depression systems often demonstrate technical feasibility through sensing, prediction, dashboards, or alerts, but these components are not always organised into a coherent system structure. The problem is not only the absence of particular technologies, but the lack of architecture-level integration that connects technical functions with safety, personalisation, clinical relevance, user engagement, and governance.

The thesis moves this state forward by deriving architecture-level requirements from multiple sources of evidence. Chapter 3 identifies fragmentation in existing IoT-based depression and emotional well-being systems. Chapters 5 and 7 then explore two design dimensions that respond to this fragmentation: contextual interpretation and engagement-oriented behavioural interaction. Chapter 9 integrates these requirements into the initial *IoTMindCare* reference architecture, while Chapter 11 refines the architecture through technical and clinical expert perspectives.

What is new is the consolidation of these previously separate concerns into a single architectural artefact. *IoTMindCare* does not simply propose another IoT application or prediction model. It organises sensing, analytics, feedback, personalisation, safety, stakeholder interaction, interoperability, and governance into a coherent reference structure. This matters because IoT-based depression management requires more than data collection or prediction; it requires a system that can interpret data, adapt to users, manage safety, support human oversight, and remain adaptable for future implementation.

The resulting state is an expert-informed reference architecture that provides a foundation for future system development. However, the architecture still needs to be implemented, piloted, and evaluated in real-world or semi-real-world settings. Future work must test whether the architectural structure can support safety, usability, personalisation, interpretability, clinical workflow integration, and future implementation in practice.

12.2.2 RQ1: Limitations and Design Requirements in Existing IoT-Based Depression Systems

RQ1 asked: *What limitations and design requirements characterise existing IoT-based approaches to depression monitoring and emotional well-being?*

In response to RQ1, the thesis argues that the main limitation of the current field is architectural fragmentation. Existing research shows that IoT-based systems can collect data, classify signals, and support some forms of emotional well-being monitoring. However, many systems remain component-focused. They often emphasise specific sensors, data sources, models, applications, or prediction tasks, while giving less attention to how these elements should be integrated with personalisation, safety, contextual interpretation, user engagement, interoperability, and governance.

Chapter 3 provides the primary evidence for this observation by systematically reviewing existing IoT-based approaches and identifying limitations across sensing, analytics, engagement, privacy, security, interoperability, and implementation-related considerations. The later chapters also confirm the importance of these limitations by addressing design areas that the review showed were underdeveloped. Chapter 5 responds to the limited role of contextual interpretation, Chapter 7 responds to the need for engagement and active user input, and Chapters 9 and 11 respond to the need for architecture-level integration and expert-informed refinement.

The contribution of RQ1 is not only descriptive. It does not simply list gaps in the literature. It establishes the design problem that the rest of the thesis addresses. The review moves the thesis beyond the observation that IoT technologies can be used for depression-related monitoring and towards the more important question of what requirements are needed for safe and personalised system design.

This matters because a field can have many technically promising prototypes while still lacking a coherent pathway towards implementation. By identifying fragmentation as the central limitation, RQ1 provides the justification for developing *IoTMindCare*. It shows why the thesis needed to move beyond isolated technical functions and towards a reference architecture that can organise the field's requirements.

What remains after RQ1 is the need to continuously update these requirements as the field develops. IoT devices, AI methods, digital health regulations, interoperability standards, and clinical workflows will continue to change. Future work should revisit the design requirements as new technologies and evidence emerge, especially in relation to explainable AI, privacy-preserving analytics, regulatory expectations, cultural appropriateness, and real-world clinical integration.

12.2.3 RQ2: Contextual and Behavioural Data for Personalised and Meaningful Monitoring

RQ2 asked: *How can the limitations identified in existing IoT-based approaches concerning environmental contextualisation and user engagement be explored to inform the design of personalised and meaningful depression monitoring and management systems?*

In response to RQ2, the thesis argues that IoT-based depression monitoring becomes more meaningful when passive sensing is combined with contextual interpretation and active user engagement. This position follows from the limitations identified in Chapter 3, where many existing systems were shown to collect behavioural, physiological, smartphone, or self-reported data without sufficiently integrating contextual or engagement-oriented mechanisms. When these data are treated as isolated signals, their interpretation can remain limited. Depression-related patterns are often ambiguous: a change in activity, sleep, phone use, or self-report may have different meanings depending on the person, their environment, their routine, and their current circumstances. Chapters 5 and 7 extend the thesis by showing how environmental context, gamified interaction, adaptive self-reporting, and behavioural indicators can support more personalised and interpretable monitoring.

The thesis advances the discussion of contextual and engagement-oriented IoT-based depression monitoring by examining contextual and behavioural dimensions as design requirements rather than as optional additions. Chapter 5 shows how temperature and humidity measurements can be processed and classified using literature-informed environmental stress-related labels, thereby motivating their inclusion as contextual variables within IoTMindCare. Chapter 7 extends this by exploring gamified interaction, adaptive self-reporting, and behavioural indicators through prototype development and synthetic-data simulation. As both studies use constructed classification labels derived from their input features, their results provide design-oriented insights. These chapters together show that IoT-based depression monitoring should not rely only on passive sensing or generic prediction. It should include mechanisms that help interpret data in context and involve the user in the monitoring process.

The contribution of RQ2 lies in positioning contextual data, behavioural interaction, and engagement within a management architecture. Environmental data are not treated as independent explanations of mental health status, and game-based behavioural tasks are not treated as standalone diagnostic tools. Instead, both are understood as partial sources of evidence that become useful when integrated with other data streams, personal baselines, and feedback mechanisms.

Meaningful monitoring requires interpretation rather than signal collection alone. Without context, a system may misread ordinary behavioural variation as risk or fail to recognise a meaningful change. Without active user interaction with the system, it may lose access to subjective experience and may not support user reflection or self-management. Through the exploration of RQ2, the thesis shows why an architecture for depression management should include both passive and active forms of data collection, and why both should be interpreted in relation to the individual.

This leaves an important need to evaluate how contextual and behavioural mechanisms perform with real users over time. Future work should examine whether environmental data improve interpretation, whether gamified interaction supports sustained engagement, whether adaptive self-reporting reduces burden, and whether users find the resulting feedback meaningful and trustworthy. Further work should also investigate how contextual and behavioural data should be weighted in relation to other data sources.

12.2.4 RQ3: Organisation of IoT Components within an Integrative Reference Architecture

RQ3 asked: *How can the literature-derived requirements and design insights from contextual sensing and engagement-oriented studies be integrated into a reference architecture that addresses safety, personalisation, stakeholder interaction, interoperability, and governance for IoT-based depression monitoring and management?*

In response to RQ3, the thesis argues that IoT-based depression management requires architecture-level organisation rather than isolated prototype functionality. The current state of the field demonstrates that many individual functions are technically possible. Sensors can collect data, algorithms can classify signals, dashboards can present information, and alerts can be generated. However, these functions do not by themselves define how a safe and personalised depression management system should operate. The missing contribution is a coherent architecture that clarifies system layers, component relationships, data flows, stakeholder roles, and cross-cutting concerns.

The thesis responds to the need for architecture-level organisation by developing *IoTMindCare*. Chapter 9 integrates requirements derived from the literature review and the contextual and behavioural studies into a reference architecture. The architecture organises sensing, communication, data processing, analytics, services, interfaces, personalisation, safety, and stakeholder interaction. It also makes explicit the cross-cutting concerns essential to mental health technologies, including privacy, governance, interoperability, and human oversight.

The architectural contribution of *IoTMindCare* is that it moves beyond identifying individual technical functions to defining how those functions should be organised within a safe, personalised, and governed system structure. This is important because depression management is not only a technical problem. It is also a clinical, personal, ethical, and organisational problem. The architecture contributes by defining how different parts of the system should relate to one another and by showing where safety, personalisation, and governance need to be embedded.

This matters because future IoT-based depression systems need a reusable structure that can guide development across different technologies and care contexts. A reference architecture provides a foundation that future systems can instantiate, evaluate, and adapt. In this sense, RQ3 represents the central architectural contribution of the thesis.

What remains is to test the architecture through future implementation research. Future work should develop prototypes based on *IoTMindCare*, examine how the layers and cross-cutting concerns operate in practice, and evaluate whether architecture-level integration improves safety, usability, interpretability, personalisation, and future implementation planning compared with component-focused systems.

12.2.5 RQ4: Expert-Informed Refinement of the Architecture

RQ4 asked: *How can technical and clinical expert perspectives inform the refinement of the proposed reference architecture to strengthen its technical and clinical relevance, safety considerations, and suitability for future implementation and validation?*

In response to RQ4, the thesis argues that an architecture for IoT-based depression management must be refined through multidisciplinary judgement before it can provide a credible basis for future implementation research. A technically coherent architecture may still be incomplete if it does not account for clinical relevance, user safety, alert burden, workflow integration, privacy, and implementation constraints. Therefore, expert-informed refinement is necessary to strengthen the feasibility, relevance, and evaluability of the architecture for future empirical work.

The thesis advances the architecture in Chapter 11, which incorporates technical and clinical expert perspectives into the refinement process. These perspectives strengthen the architecture by clarifying the need for personalised baselines, adaptive calibration, multimodal fusion, alert governance, human oversight, hybrid edge-cloud processing, privacy, interoperability, and future implementation considerations. This refinement builds on the initial architecture presented in Chapter 9 and transforms it into a more clinically and technically informed reference structure.

Expert input is incorporated into the architecture refinement process rather than being treated only as general feedback. The refinement identifies where the initial architecture lacked sufficient operational detail, particularly regarding calibration, escalation, safety mechanisms, and edge-cloud distribution. It also reinforces the importance of human-in-the-loop support, showing that automated outputs must be governed rather than treated as independent clinical decisions.

The relevance of expert-informed refinement is especially important because IoT-based depression systems involve safety-sensitive decisions. Frequent false alerts may increase user burden, reduce trust, and create unnecessary escalation, while missed signs of meaningful deterioration may delay support. Clinical experts can identify risks in interpretation, escalation, and workflow, while technical experts can identify risks in infrastructure, interoperability, data processing, and scalability. Therefore, RQ4 strengthens the thesis by ensuring that *IoTMindCare* is not only conceptually integrated, but also more feasible and relevant for future implementation.

The expert findings revealed variation in perspectives and highlighted unresolved technical and clinical considerations. These included challenges related to missing-data handling, edge-cloud processing, personalisation, alert reliability, and clinical oversight. The feedback highlighted the need to balance competing priorities, particularly reducing false alerts while minimising missed critical events. These differences informed the architectural refinements while identifying areas requiring further investigation and evaluation.

The next stage is to move from expert-informed refinement to empirical evaluation. Future work should evaluate the refined architecture with users, caregivers, clinicians, and healthcare organisations. This includes testing alert governance, user control, caregiver boundaries, clinician workload, privacy mechanisms, edge-cloud processing, and interoperability pathways in real or semi-real pilot contexts. Only through such evaluation can the architecture move beyond expert-informed design towards evidence-based implementation research.

12.3 Synthesis of Thesis-Level Findings

This section synthesises the broader thesis-level findings that cut across the publication-based chapters. These findings interpret the significance and value of the research contributions and how this research moves the field towards architecture-level integration. The synthesis draws together the contributions developed across the systematic literature review, the contextual and behavioural studies, the *IoTMindCare* reference architecture, and its expert-informed refinement. It is organised around six transitions that show how the thesis moves from fragmented IoT-based monitoring towards a safe, personalised, context-aware, governed, and implementation-oriented architecture for depression management.

Figure 12.1 summarises these six thesis-level transitions before they are discussed in detail. The figure shows how the thesis moves from fragmented or limited approaches in the existing literature towards the architecture-level contribution developed through *IoTMindCare*.

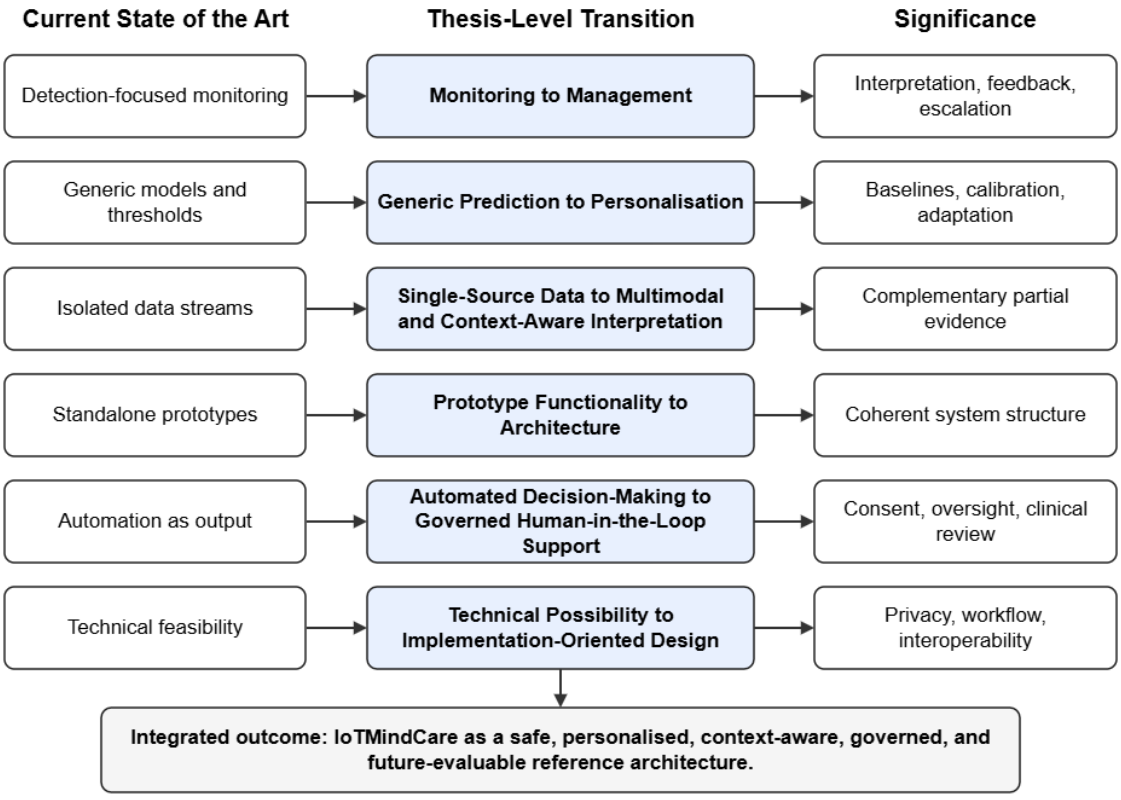


Fig. 12.1 The six key transitions offered through *IoTMindCare*.

12.3.1 From Monitoring to Management

Existing IoT-based depression and emotional well-being research has shown that continuous or near-continuous monitoring is technically feasible. Smartphones, wearables, environmental sensors, self-report tools, and machine learning models can collect and analyse patterns that may be difficult to capture through periodic clinical encounters or retrospective self-report alone [27]. However, this capability does not automatically amount to depression management. Monitoring identifies possible changes; management requires those changes to be interpreted, contextualised, communicated, and acted upon safely.

This distinction is important because depression is not a single event that can be addressed through a single detection point. It is a fluctuating condition shaped by individual experience, daily routines, behaviour, sleep, environment, social context, and help-seeking patterns [14]. Therefore, an IoT-based system designed to support people with depression needs to operate over time. It must not only identify patterns but also interpret them in context, adapt to the individual, support reflection, and determine when information should remain with the user, when it may be appropriate for a caregiver, and when it may require clinical attention.

Chapter 3 showed that much of the literature remains closer to monitoring than management, with many systems focusing on data acquisition, classification, or prediction, but giving less attention to what happens after a signal is detected [20]. Chapter 5 demonstrates that contextual information can help explain why behavioural or emotional patterns may occur [24], while Chapter 7 shows that user interaction, adaptive self-reporting, and feedback can support engagement and reflection [25]. These elements are then integrated in *IoTMindCare* in Chapter 9, where sensing, analytics, feedback, personalisation, safety, and stakeholder interaction are organised as connected parts of a management pathway [26].

The contribution is the repositioning of monitoring within a broader management architecture. This moves the thesis beyond the question of whether IoT systems can detect depression-related signals and towards how such signals can support safe, adaptive, and meaningful action. The expert-informed architectural refinement in Chapter 11 strengthens this position by clarifying that management requires calibration, alert governance, escalation pathways, and oversight. These are not secondary implementation details; they are core architectural requirements for using monitoring outputs responsibly.

The integration of contextual interpretation, user interaction, architectural organisation, and expert-informed refinement reframes sensing and prediction as parts of a broader management pathway rather than as isolated technical outputs. In this pathway, data collection is connected to interpretation, feedback, personalisation, and governed escalation. A technically accurate signal has limited value if it is not connected to a meaningful support process.

The next step is to evaluate this management pathway in practice. Future work should implement the architecture and examine whether the proposed flow from sensing to interpretation, feedback, and escalation can operate safely and usefully over time. Particular attention should be given to user responses to feedback, caregiver and clinician interpretation of system outputs, alert governance, and the balance between continuous monitoring, privacy, burden, and trust.

12.3.2 From Generic Prediction to Personalisation

Many IoT-based depression systems use population-level models, fixed features, or general prediction approaches. These approaches are useful because they show that mental health-related signals can be modelled using behavioural, physiological, smartphone, or self-reported data. However, depression varies substantially between individuals. Symptoms, severity, triggers, routines, coping strategies, engagement capacity, sleep patterns, social behaviour, and responses to environmental or behavioural changes are not uniform across individuals. The same signal may have different meanings for different users.

Individual variation creates a limitation for systems that rely mainly on generic models or fixed thresholds. Reduced activity, increased phone use, disrupted sleep, or reduced interaction may indicate deterioration for one user, but normal variation or situational change for another. If these signals are interpreted without reference to the individual's own baseline and context, the system may over-alert some users, under-detect changes in others, or provide feedback that feels irrelevant. These failures are not only technical; they affect trust, engagement, safety, and willingness to continue using the system.

The thesis responds to the limitations of generic interpretation by treating personalisation as an architectural requirement rather than an optional enhancement. Chapter 3 identified insufficient personalisation as a recurring limitation in the field [20]. The contextual and behavioural studies in Chapters 5 and 7 then show why universal interpretation is insufficient: environmental data are meaningful only in relation to the user's setting, while behavioural interaction and self-reporting are influenced by motivation, fatigue, cognitive load, preference, and symptom state [24, 25]. *IoTMindCare* incorporates personalisation as a cross-cutting concern, and expert-informed refinement makes this more explicit through personalised baselines, calibration, adaptive thresholds, progressive learning, and model-drift awareness.

The contribution is not simply that the thesis recommends personalised feedback. It advances a broader view in which personalisation affects sensing, interpretation, analytics, feedback, alerting, and safety. A personalised architecture can begin with general models where needed, but it must progressively adapt as user-specific data becomes available. This

changes the role of prediction: prediction is no longer a standalone output, but part of an ongoing calibration process.

The combined contribution creates a new state in which personalisation is understood as a safety-sensitive architectural function. A system that is not personalised may misinterpret normal variation as risk, or risk as normal variation. By contrast, a personalised system can interpret signals in relation to the user's history, context, and changing patterns. This does not remove uncertainty, but it creates a more meaningful basis for support.

Future research needs to test how this personalisation process should operate in practice. This includes determining how long it takes to establish reliable baselines, how adaptive thresholds should be updated, how to avoid normalising gradual deterioration, and how personalised models can remain explainable and auditable. Further work is also needed to balance personalisation with privacy, transparency, and clinical oversight.

12.3.3 From Single-Source Data to Multimodal and Context-Aware Interpretation

Existing research shows that different data sources can contribute to depression-related monitoring. Passive sensor data can capture changes in movement, routine, sleep, location, or device use. Self-report can capture subjective experience. Behavioural tasks can provide structured indicators of attention, reaction, or cognitive performance. Environmental data can describe contextual conditions. Each of these sources is useful, but each is also partial. Depression affects and is affected by multiple dimensions of life, including mood, cognition, behaviour, activity, sleep, environment, social interaction, and self-perception [234]. No single data stream can fully represent this complexity.

The limitation of single-source interpretation is ambiguity. Passive sensor data may show that behaviour has changed, but not why. Self-report may provide direct insight, but may be affected by low motivation, recall bias, or inconsistent engagement [15]. Game-based tasks may provide behavioural indicators, but performance may be affected by fatigue, familiarity, attention, or willingness to participate. Environmental data can provide context, but cannot independently explain mental health status. Therefore, each modality provides partial evidence, and the meaning of that evidence depends on how it relates to other signals.

The thesis responds to the ambiguity of single-source interpretation by arguing that multimodal data require architectural interpretation rather than simple accumulation. Chapter 3 identified fragmentation in sensing approaches and limited integration across data types [20]. Chapter 5 contributes an environmental interpretive layer, while Chapter 7 combines behavioural interaction, adaptive self-reporting, sensor input, and environmental

context [24, 25]. These strands are then brought together in *IoTMindCare* in Chapter 9, which provides a structure for organising heterogeneous data streams and connecting them to analytics, feedback, personalisation, and safety mechanisms [26].

The contribution is a shift from collecting more data to interpreting relationships among data. More data does not automatically produce a better understanding. Without a mechanism to weigh, prioritise, contextualise, and govern different signals, multimodal systems may become more complex without becoming more clinically meaningful. The expert-informed architectural refinement in Chapter 11 strengthens this contribution by emphasising multimodal fusion, calibration, context-aware interpretation, and alert governance.

The combined outcome is a context-aware model of multimodal integration. In this model, environmental, behavioural, self-reported, and sensor-based data are not treated as independent inputs that automatically produce insight. They are complementary forms of partial evidence. Meaning emerges from their relationships with each other, the user's baseline, the context in which the data are produced, and the action pathway associated with the interpretation.

Further work should examine how this multimodal interpretation should be implemented and evaluated. Future studies need to investigate how data streams should be weighted, how missing or unreliable data should be handled, how uncertainty should be communicated, and which combinations of signals are most useful for user reflection, caregiver awareness, or clinical escalation. Longitudinal evaluation is needed to determine whether context-aware multimodal interpretation improves trust, usefulness, and decision support compared with single-source or non-contextual approaches.

12.3.4 From Prototype Functionality to Architecture

The field has demonstrated many useful prototype functions. Sensors can collect data, applications can gather self-reports, algorithms can classify signals, dashboards can display outputs, and alerts can be generated [235, 27]. These functions show technical feasibility, but they do not by themselves solve the system-level problem of safe and personalised depression management. What is often missing is a clear account of how these functions should fit together.

This thesis responds by moving from isolated functionality to architecture-level integration. The systematic review in Chapter 3 established that the field is not lacking technical possibilities; rather, it lacks sufficient integration across sensing, analytics, feedback, personalisation, safety, privacy, stakeholder interaction, interoperability, and governance [20]. The contextual and behavioural studies in Chapters 5 and 7 then provide design evidence for two important architectural dimensions: contextual interpretation and user engagement [24, 25].

These requirements are integrated in Chapter 9 through *IoTMindCare*, and refined in Chapter 11 through technical and clinical expert perspectives.

The contribution of *IoTMindCare* in Chapter 9 is that it provides a reusable reference structure rather than another single-purpose prototype. It identifies the major layers required for IoT-based depression management, including sensing, communication, data processing, services, interfaces, personalisation, and safety. It also clarifies relationships between components and makes cross-cutting concerns such as privacy, governance, interoperability, and stakeholder roles explicit. This is important because safe depression management depends not only on whether each function works individually, but on whether the functions operate coherently as a system.

The thesis provides an architecture-level foundation for future IoT-based depression management systems. In this architecture, contextual sensing, behavioural interaction, self-reporting, analytics, feedback, stakeholder roles, safety mechanisms, and governance are no longer treated as separate concerns. They are organised into one coherent architectural artefact. This changes how the thesis contributions fit together: Chapters 3, 5, and 7 provide design dimensions; Chapter 9 integrates them; and the expert-informed refinement in Chapter 11 strengthens their feasibility and relevance.

This architectural contribution needs evidence of implementation. A reference architecture provides structure, but it does not determine every implementation decision. Future work should instantiate *IoTMindCare* in prototypes, test different sensor and analytics configurations, evaluate interface designs, and examine how the architecture can be adapted to different user groups, clinical workflows, and technology environments. Comparative studies could also examine whether systems based on the architecture better support safety, interpretability, usability, and future implementation planning than systems without architecture-level integration.

12.3.5 From Automated Decision-Making to Governed Human-in-the-Loop Support

Automated methods are increasingly used in digital mental health and IoT-based monitoring for pattern recognition, prediction, recommendation, and alert generation. These methods can support earlier awareness and more continuous monitoring. However, automation becomes problematic if it is treated as a substitute for interpretation, judgement, or care. Depression management involves ambiguity, personal context, risk, changing circumstances, and clinical judgement. Therefore, automated outputs need governance.

This issue is especially important in mental health because the consequences of system error are not limited to technical inconvenience. False positives may create anxiety, unnecessary escalation, or alert fatigue. False negatives may delay support. Poorly designed feedback may make users feel blamed, monitored, or misunderstood. Inappropriate data sharing may damage privacy, autonomy, and trust. For this reason, safety cannot be treated as a minor technical feature.

The thesis contributes a governed human-in-the-loop model of IoT-based depression support. Chapter 3 identified safety, privacy, implementation, and governance limitations in existing systems [20]. *IoTMindCare* in Chapter 9 incorporates safety and stakeholder interaction as architectural concerns, while the expert-informed refinement operationalises these concerns through persistence windows, cooldown periods, escalation logic, alert prioritisation, and clinical review [26]. These mechanisms clarify that automated systems may assist with monitoring and early warning, but should not independently determine clinical meaning or replace human judgement.

The governed human-in-the-loop model reframes automation as a form of supported decision-making rather than as an independent technical capability. The architecture does not reject automation; it structures its role. The system can identify patterns, support reflection, and assist in timely response, but escalation must be governed by rules, context, consent, and human oversight. This also recognises the different roles of users, caregivers, and clinicians. Users need control and transparency. Caregiver access may be useful, but must be bounded. Clinicians need interpretable and prioritised information rather than unfiltered data or excessive alerts.

A safety-oriented and human-centred architecture requires human involvement to be built into the system rather than added only after automated outputs are generated. In *IoTMindCare*, human involvement is supported through stakeholder roles, alert governance, safety mechanisms, and clinical review pathways. This makes the architecture more appropriate for mental health, where data are sensitive, and interpretation requires judgement.

Future work should evaluate how this governance model operates in practice. This includes examining how alerts should be prioritised, how persistence windows and cooldown mechanisms affect safety and burden, how clinicians and caregivers interpret outputs, and how users understand and control escalation pathways. Further research should also examine how human-in-the-loop processes can be integrated into clinical workflows without creating excessive workload or unclear responsibility.

12.3.6 From Technical Possibility to Implementation-Oriented Design

IoT-based depression research has demonstrated that many technical components are possible. Data can be collected from devices, processed using machine learning models, displayed through dashboards, and used to generate feedback or alerts. However, technical possibility does not mean that a system is ready for implementation, piloting, or clinical use. Future implementation would introduce requirements that are often underdeveloped in prototype studies, including reliability, privacy, user acceptance, clinician workload, interoperability, regulatory awareness, data governance, device heterogeneity, and long-term maintenance.

The thesis responds to the gap between technical feasibility and future implementation needs by incorporating implementation-oriented considerations into the architecture. Chapter 3 identified implementation-related limitations, interoperability, privacy, and governance as recurring issues in existing work [20]. These concerns are organised within the initial *IoTMindCare* architecture in Chapter 9 and strengthened through expert-informed refinement in Chapter 11, particularly in relation to hybrid edge-cloud processing, alert governance, clinical oversight, privacy, interoperability, and implementation-oriented design. Chapters 5 and 7 also contribute by showing that future systems must account for real user contexts, engagement patterns, and the interpretive limits of data [24, 25].

The implementation-oriented contribution of the thesis is the movement from asking whether IoT-based depression support can be technically built to identifying the architectural conditions that should be considered before such systems are responsibly implemented or evaluated. The edge-cloud processing view illustrates the importance of these architectural conditions. A cloud-based system may support scalable analytics and longitudinal learning but may also raise concerns about latency, connectivity, privacy, and resilience. An edge-based system may support local responsiveness and reduce data transfer, but may limit long-term learning, model retraining, and integration with broader services. The refined architecture in Chapter 11 supports a hybrid design view in which local processing may be used for time-sensitive and privacy-sensitive functions, while cloud infrastructure may support longitudinal analytics, model updating, storage, and interoperability.

The thesis identifies implementation-related requirements that should be addressed before future piloting, evaluation, or possible clinical integration can be considered. In this architecture, privacy, governance, interoperability, clinical workflow, edge-cloud processing, and user acceptance are treated as design concerns rather than later-stage additions. This matters because systems that ignore these concerns during design may become difficult, unsafe, or unacceptable to implement later.

Future work should instantiate and evaluate *IoTMindCare* through staged prototype development and piloting. This should include technical testing of edge-cloud allocation,

evaluation of interoperability pathways, privacy and consent mechanisms, clinical workflow considerations, user acceptance, caregiver role boundaries, and long-term system maintenance. Further work should also consider regulatory classification, cultural appropriateness, equity, and organisational adoption.

12.4 Overall Contribution of the Thesis

The overall contribution of this thesis is the development and expert-informed refinement of *IoTMindCare*, an integrative reference architecture for safe and personalised IoT-based depression monitoring and management. The significance of this contribution lies in its synthesis. The thesis does not simply add another sensing model, mobile application, or machine learning experiment to the literature. Instead, it addresses the more fundamental issue of how sensing, interpretation, personalisation, engagement, safety, stakeholder interaction, interoperability, and governance can be organised into a coherent system.

The thesis contribution is supported by architectural modelling, ATAM-informed analysis, experimental and simulation-based studies, and expert-informed refinement. Together, these provide evidence for the proposed design, its technical considerations, and clinical relevance, while establishing a foundation for further evaluation of operational performance and clinical outcomes.

Many of the challenges in IoT-based depression management are interdependent. Technical feasibility, clinical relevance, personalisation, safety, and future implementation considerations cannot be treated as separate concerns. Decisions made in one area affect the others. For example, personalisation affects alert reliability; alert reliability affects clinical trust; privacy affects user engagement; and interoperability affects whether system outputs can be integrated into care. The contribution of *IoTMindCare* is that it provides a structure for considering these dependencies together.

At the conceptual level, the thesis reframes IoT-based depression support as a socio-technical management problem. This reframing recognises that depression management involves more than data collection and prediction. It also requires interpretation, personalisation, user engagement, safety, human oversight, privacy, governance, and clinical relevance.

At the architectural level, the thesis contributes *IoTMindCare* as a reusable reference structure that can guide future system development. The architecture identifies the core layers and cross-cutting concerns needed for safe and personalised IoT-based depression management, while remaining adaptable to different technologies, populations, and care contexts.

At the methodological level, the thesis demonstrates a design-oriented pathway from problem identification to architecture development and expert-informed refinement. The contribution of this thesis lies in the cumulative argument formed across the research process: identifying fragmentation in existing IoT-based depression systems, exploring contextual and behavioural design requirements, developing the initial *IoTMindCare* reference architecture, and refining it through technical and clinical expert perspectives. This progression is developed from the systematic review in Chapter 3, through the contextual and behavioural studies in Chapters 5 and 7, to the architecture development and expert-informed refinement in Chapters 9 and 11.

At the practical level, the thesis identifies design requirements for future systems, including multimodal sensing, contextual interpretation, personalised calibration, adaptive engagement, alert governance, hybrid edge-cloud processing, privacy protection, interoperability, and clinical oversight. These requirements provide a foundation for future implementation research, piloting, and empirical evaluation.

Taken together, the conceptual, architectural, methodological, and practical contributions show the integrative nature of the thesis. *IoTMindCare* brings together technical, clinical, behavioural, contextual, and governance considerations within one architectural framework. The significance of this contribution is that it provides a structured way to move IoT-based depression research beyond isolated prototypes and prediction models towards safer, personalised, context-aware, and future-evaluable systems. This matters because IoT-based depression systems may be technically feasible but still unsafe, intrusive, difficult to trust, or difficult to implement if personalisation, safety, stakeholder roles, privacy, interoperability, and governance are not considered together. The thesis provides an architecture-level foundation for responsible future implementation, while remaining appropriately bounded because it does not claim clinical effectiveness or deployment validation.

12.5 Limitations and Threats to Validity

The individual publication-based chapters have already discussed limitations specific to each study. Therefore, this section focuses on limitations that apply across the thesis as a whole. To structure this discussion, an adapted threats-to-validity framework is used, drawing on validity categories commonly used in empirical software engineering: construct validity, internal validity, external validity, conclusion validity, and reliability [236]. The framework is adapted because this thesis is not a single controlled experiment, but a thesis by publication that combines systematic review, design-oriented studies, reference architecture development, and expert-informed refinement.

Table 12.1 summarises the main thesis-level threats to validity and identifies the chapters most relevant to each category. The table provides an overview; the following subsections explain how these threats bound the contribution of the thesis.

Table 12.1 Summary of thesis-level threats to validity.

Validity category	**Main threat in this thesis**	**Relevant chapters**
Construct validity	The key constructs, including safety, personalisation, context-awareness, and future-adaptability, are broad and multidimensional. The thesis examines these constructs at the architecture level rather than through full clinical deployment.	Chapters 3–11
Internal validity	The thesis supports architecture-level findings, but does not provide causal evidence that *IoTMindCare* improves depression outcomes, diagnostic accuracy, safety, or clinical effectiveness in practice.	Chapters 9 and 11
External validity / transferability	*IoTMindCare* is designed to be adaptable, but its transfer to different users, healthcare systems, technologies, cultural settings, and clinical workflows has not yet been tested.	Chapters 9 and 11
Conclusion validity	There is a risk of overclaiming because the thesis includes multiple published studies and expert-informed refinement, but not full deployment or clinical validation.	Chapters 3–11
Reliability / dependability	Architecture development involves interpretive and design-oriented judgement, particularly in organising layers, components, stakeholder roles, and cross-cutting concerns.	Chapters 3–11

12.5.1 Construct Validity

Construct validity concerns whether the thesis has adequately captured the concepts it claims to address. In this thesis, the key constructs include safe, personalised, context-aware, and future-adaptable IoT-based depression monitoring and management. These constructs are broad and include clinical, technical, ethical, behavioural, and organisational dimensions.

The thesis addresses this challenge through its cumulative structure. Chapter 3 identifies the range of limitations and design requirements in existing IoT-based depression and emotional well-being systems. Chapters 5 and 7 examine selected dimensions of contextual interpretation and engagement. Chapter 9 integrates these dimensions into the initial *IoTMindCare* architecture, and Chapter 11 refines the architecture through technical and clinical expert perspectives.

However, the thesis examines these constructs at the architecture level rather than through clinical deployment or outcome evaluation. For example, it explains how safety and personalisation should be organised within *IoTMindCare*, but it does not demonstrate that these mechanisms improve clinical outcomes in practice. Similarly, environmental variables such as temperature and humidity support contextual interpretation, but they do not represent the full social, cultural, interpersonal, or clinical context of depression management. The architecture should be interpreted as a structured foundation for future development, not as a complete representation of all factors involved in depression care.

12.5.2 Internal Validity

Internal validity concerns the strength of the relationship between the evidence developed in the thesis and the findings from that evidence. The central finding of the thesis is that fragmented IoT-based depression monitoring and management research can be consolidated into an integrative reference architecture. The evidence developed across the thesis supports this architecture-level finding.

The main limitation is that the thesis does not provide causal evidence that *IoTMindCare* improves depression management, reduces symptoms, improves safety, or enhances clinical outcomes. The findings support claims about architectural organisation, design requirements, expert-informed refinement, and future implementation considerations. They do not support claims of clinical effectiveness or diagnostic validity.

The evidence developed in the thesis supports architecture-level claims rather than causal or clinical-effectiveness claims. Therefore, *IoTMindCare* is positioned as a reference architecture rather than as a completed clinical intervention. The expert-informed refinement in Chapter 11 strengthens the feasibility and relevance of the architecture, but it does not replace empirical testing with users, caregivers, clinicians, and healthcare organisations. Further implementation and evaluation are required before stronger claims can be made about real-world impact. As the architectural evaluation draws on predefined requirements and scenarios, independent evaluation would provide additional evidence beyond the demonstrated alignment between design decisions and quality attributes.

12.5.3 External Validity and Transferability

External validity concerns the extent to which the thesis findings can be transferred to other settings, populations, technologies, or healthcare contexts. *IoTMindCare* is intentionally designed as a reference architecture, which supports adaptability across different sensors, platforms, analytics methods, user groups, and care models.

The adaptability of *IoTMindCare* does not imply direct transferability without modification. Different contexts may require different sensor configurations, escalation pathways, privacy settings, governance arrangements, cultural considerations, and clinical workflows. For example, a system designed for independent adults living at home may require different caregiver involvement and safety mechanisms from a system designed for students, older adults, or people receiving formal clinical care. Healthcare systems may also differ in interoperability standards, regulatory requirements, workload pressures, and digital infrastructure.

The thesis supports transferability at the architectural level, but not universal applicability at the implementation level. Future work must instantiate and evaluate *IoTMindCare* in specific contexts to determine how the architecture should be adapted and whether its components operate effectively in practice.

12.5.4 Conclusion Validity

Conclusion validity concerns whether the conclusions drawn from the thesis are proportionate to the evidence. The main risk is overclaiming. Because the thesis includes multiple published studies and expert-informed refinement, it could be tempting to present *IoTMindCare* as deployment-ready or clinically proven. This thesis avoids that claim.

The conclusion is deliberately bounded: *IoTMindCare* provides a coherent and expert-informed reference architecture for organising IoT-based depression monitoring and management requirements. It does not provide evidence that the architecture improves depression outcomes, achieves diagnostic validity, or is ready for direct clinical deployment.

The bounded scope of the conclusion is important because the field is also rapidly evolving. Developments in IoT, AI, interoperability standards, privacy-preserving analytics, regulation, and digital mental health practice may introduce new requirements. The conclusions should be interpreted as an architecture-level contribution at this stage of the field, while recognising that future updates and empirical validation will be necessary.

12.5.5 Reliability and Dependability

Reliability concerns whether the research process is transparent and dependable. In this thesis, dependability is supported by the publication-based structure, the systematic review, the explicit links between chapters, and the expert-informed refinement of the architecture.

The reliability of the research process must also be considered in relation to the interpretive and design-oriented nature of architecture development. Decisions about layers, components, stakeholder roles, safety mechanisms, and cross-cutting concerns are based on

synthesis rather than mechanical derivation. This introduces a degree of researcher interpretation. The thesis addresses this by making the research pathway explicit: fragmentation is identified through the review, contextual and behavioural dimensions are explored in the design-oriented chapters, the architecture is developed from these requirements, and expert-informed refinement is used to strengthen the design.

Future implementation across multiple contexts would further test the dependability of the architecture. Such work would show whether the architectural structure remains stable across different settings or whether additional components, layers, or governance mechanisms are required.

12.6 Future Work

This thesis establishes *IoTMindCare* as an architecture-level foundation for safe and personalised IoT-based depression monitoring and management. The next stage of research is to examine how this foundation can be instantiated, tested, adapted, and evaluated in practice. Future work should move progressively from technical implementation to stakeholder evaluation, longitudinal piloting, and clinical and organisational validation.

12.6.1 Prototype Implementation of IoTMindCare

A future direction is to implement *IoTMindCare* as a modular prototype. The architecture developed in Chapter 9 and refined in Chapter 11 defines the main layers, components, and cross-cutting concerns required for safe and personalised IoT-based depression management. However, implementation is needed to examine how these components operate together in practice.

A future project could develop a working prototype that includes sensing, data ingestion, preprocessing, analytics, feedback, personalisation, alert governance, and stakeholder interfaces. The prototype should be modular so that different sensors, machine learning models, feedback mechanisms, and interface designs can be substituted and evaluated. This would allow *IoTMindCare* to be tested as a flexible reference architecture rather than as a single fixed application.

This work would address several practical questions:

- How should passive sensing, environmental data, behavioural interaction, and self-report be integrated?
- How should responsibilities be distributed between edge and cloud components?

- How should personalised baselines and adaptive thresholds be represented technically?
- How should outputs be communicated differently to users, caregivers, and clinicians?

Addressing these questions would move *IoTMindCare* from architectural design towards technical feasibility.

12.6.2 Simulation-Based Evaluation of Safety and Alert Governance

Future work should also evaluate safety-related mechanisms through simulation before any real-world piloting or user-facing evaluation. The thesis has argued that depression management requires more than detection; it requires calibration, escalation logic, alert governance, and human oversight. These mechanisms should be tested under controlled conditions before being used with real users.

A future project could develop simulation scenarios representing different user profiles, behavioural routines, environmental conditions, self-report patterns, and risk trajectories. These scenarios could be used to test how the system responds to stable patterns, gradual deterioration, sudden changes, missing data, conflicting signals, false positives, false negatives, sensor failure, and connectivity loss. This would extend the safety and personalisation work developed in Chapters 9 and 11.

This simulation would help refine key design decisions. For example:

- How long should a concerning pattern persist before escalation occurs?
- How should the system respond when passive sensing and self-report conflict?
- How can cooldown periods reduce alert fatigue without suppressing meaningful deterioration?
- How should the system behave when cloud connectivity is unavailable?

Simulation would allow these questions to be examined before exposing users to higher-risk deployment conditions.

12.6.3 User-Centred Evaluation with Key Stakeholders

Another future direction is to evaluate the system with the stakeholders who would be affected by it. Chapter 7 highlighted the importance of engagement, adaptive self-reporting, and feedback, while Chapter 11 highlighted clinical relevance, safety, and stakeholder interaction.

Future work should examine how users, clinicians, caregivers, and technical stakeholders understand and experience the system.

A future project could involve people with lived experience of depression, clinicians, and caregivers in participatory design, usability testing, and acceptability evaluation. The study could examine whether the system is understandable, trustworthy, manageable, and acceptable. Particular attention should be given to feedback design, privacy preferences, user burden, consent, caregiver access, alert interpretation, and clinician workflow.

This work is necessary because a technically coherent architecture may still fail if users experience it as intrusive, burdensome, or unclear. Similarly, clinicians may reject a system if it generates excessive alerts or provides outputs that do not fit clinical workflows. Caregiver involvement may be useful, but it must be carefully bound by consent, user autonomy, and safety. This future work would strengthen the human-centred and ethical foundations of *IoTMindCare*.

12.6.4 Longitudinal Pilot Evaluation

A future direction is longitudinal pilot evaluation in real-world or semi-real-world settings. Depression-related monitoring depends on patterns over time, and many of the key architectural features of *IoTMindCare*, including personalised baselines, adaptive thresholds, contextual interpretation, engagement, and alert governance, can only be properly examined longitudinally.

A future pilot study could evaluate whether *IoTMindCare* can collect reliable multimodal data over several weeks or months, support personalised baselines, generate meaningful feedback, and manage alerts safely. The study should examine data quality, missing data, engagement over time, user burden, technical reliability, privacy concerns, and the usefulness of feedback. It should also examine whether contextual and behavioural data improve interpretation compared with single-source or non-contextual monitoring.

This work would extend the environmental and behavioural studies in Chapters 5 and 7. Those chapters demonstrated the potential value of contextual sensing and engagement-oriented design. A longitudinal pilot would test whether these design dimensions remain useful in everyday use and whether they can be integrated into a safe and acceptable monitoring pathway.

At this stage, the focus should remain on feasibility, safety, and acceptability rather than definitive clinical effectiveness. The expected contribution would be early feasibility evidence showing whether the architecture can operate reliably and meaningfully over time in a pilot context.

12.6.5 Clinical and Organisational Validation

If feasibility, safety, and acceptability are established, future work should examine whether *IoTMindCare*-based systems can support clinically meaningful outcomes and fit within real care pathways. This would represent a later stage of research, after technical implementation, simulation, stakeholder evaluation, and pilot testing.

A future study could examine whether the system supports earlier identification of deterioration, improved self-management, better communication between users and care providers, improved continuity of care, or more timely escalation. Such work would require appropriate ethical approval, clinical governance, risk management, and carefully selected outcome measures. It may also require comparison with existing digital monitoring or care approaches.

Clinical validation should be accompanied by organisational evaluation. A system may be clinically promising but still difficult to adopt if it increases clinician workload, produces unclear responsibilities, lacks interoperability, or creates governance concerns. Future research should examine alert management, workflow integration, interoperability with health information systems, regulatory classification, data governance, service models, and long-term maintenance.

12.6.6 Cross-Cutting Future Research Directions

Several cross-cutting research directions should be pursued alongside the staged implementation pathway:

1. Privacy-preserving analytics should be investigated because *IoTMindCare* may involve sensitive behavioural, environmental, and self-reported data. Future work could examine edge-based processing, data minimisation, federated learning, consent-aware sharing, and secure audit mechanisms.

2. Explainable AI should be explored so that users, caregivers, and clinicians can understand why feedback or alerts are generated. Explainability is particularly important in mental health because unclear outputs may reduce trust, increase anxiety, or limit clinical usefulness.

3. Interoperability requires further development. Future work should examine how *IoTMindCare* outputs can be integrated with clinical dashboards, electronic health records, digital mental health platforms, or other health information systems without overwhelming clinicians or compromising privacy.

4. Cultural and equity considerations should be examined more deeply. Depression management, help-seeking, family involvement, privacy expectations, and technology use may vary across cultural and social groups. Future research should investigate how *IoTMindCare* can be adapted in culturally safe and equitable ways.

5. Long-term governance models need further investigation. Future work should clarify who is responsible for monitoring alerts, updating models, managing thresholds, maintaining devices, auditing system outputs, responding to risk, and managing consent over time. These governance questions are essential if IoT-based depression management is to move from research design towards responsible practice.

12.7 Chapter Conclusion

This chapter has discussed the thesis as an integrated body of work. The chapter first addressed the research questions and then synthesised six thesis-level findings that explain the broader significance of the contribution:

- From monitoring to management.
- From generic prediction to personalisation.
- From single-source data to multimodal and context-aware interpretation.
- From prototype functionality to architecture.
- From automated decision-making to governed human-in-the-loop support.
- From technical possibility to implementation-oriented design.

These findings show that the contribution of the thesis is not limited to the individual publications. The main contribution is the development and expert-informed refinement of *IoTMindCare* as an integrative reference architecture. The architecture addresses fragmentation in the field by connecting sensing, environmental context, behavioural engagement, self-reporting, analytics, personalisation, feedback, safety, human oversight, privacy, interoperability, and future implementation considerations within one coherent framework.

The thesis contributes to IoT, digital mental health, and system architecture research by showing how safe and personalised depression management can be approached as an integrated socio-technical design problem. *IoTMindCare* should not be interpreted as a finished clinical intervention or a replacement for clinical care. Rather, it provides a structured

and expert-informed foundation for future implementation, piloting, and evaluation of IoT-based depression management systems. The significance of this contribution is that it provides a bridge between technical feasibility and responsible future implementation by showing how IoT-based depression systems can be designed with integration, personalisation, safety, governance, and human oversight in mind from the outset.

Chapter 13

Conclusion

13.1 Introduction

This thesis addressed the architectural fragmentation of current IoT-based approaches to depression monitoring and management. Existing research has shown that IoT devices, mobile applications, environmental sensors, self-report tools, machine learning models, dashboards, and alerts can each contribute to depression-related monitoring. However, these components are often developed as isolated technical functions. The central concern of this thesis was not simply whether IoT technologies can collect data or generate predictions, but how these technologies can be organised into a safe, personalised, context-aware, and clinically meaningful system.

This research responded to this problem by developing and refining *IoTMindCare*, an expert-informed reference architecture for IoT-based depression monitoring and management. The contribution of the thesis is architecture-level rather than intervention-level. The architecture provides a structured foundation for organising the sensing, interpretation, engagement, personalisation, safety, stakeholder interaction, interoperability, and governance requirements that future IoT-based depression systems need to address.

13.2 Summary of the Thesis

The thesis developed its contribution progressively. The systematic literature review in Chapter 3 first established that IoT-based depression and emotional well-being research has made technical progress but remains fragmented. Existing studies demonstrate the potential of sensors, applications, algorithms, and monitoring tools, but they often give limited attention

to how these components should work together within a broader management system. This finding established the need for an architecture-level response.

After identifying the fragmentation, this research examined two important design dimensions that informed the development of the architecture: environmental contextualisation and behavioural engagement. The environmental study in Chapter 5 showed that contextual data can support more meaningful interpretation of behavioural and emotional patterns. Environmental information does not independently explain mental health status, but it can help clarify the conditions in which changes occur. The behavioural and gamified study in Chapter 7 extended the thesis beyond passive sensing by showing the importance of user interaction, adaptive self-reporting, and engagement-oriented design. Together, these studies showed that IoT-based depression management requires both contextual interpretation and active user participation.

The contextual and behavioural design requirements identified through Chapters 5 and 7 were then integrated into the initial *IoTMindCare* reference architecture, presented in Chapter 9. The architecture organised the main layers and cross-cutting concerns needed for IoT-based depression management, including sensing, communication, data processing, analytics, services, interfaces, personalisation, safety, stakeholder interaction, interoperability, and governance. The central architectural contribution of the thesis is the movement from isolated technical components towards a coherent system-level structure.

The final architectural refinement stage in Chapter 11 strengthened the architecture through expert-informed evaluation. Technical and clinical expert perspectives highlighted the importance of personalised baselines, adaptive calibration, multimodal fusion, alert governance, human oversight, privacy, hybrid edge-cloud processing, interoperability, and future implementation planning. This refinement helped ensure that *IoTMindCare* was not only technically structured, but also more relevant to the practical and safety-sensitive realities of future IoT-based mental health systems.

13.3 Overall Contribution

The main contribution of this thesis is the development and expert-informed refinement of *IoTMindCare* as an integrative reference architecture for safe and personalised IoT-based depression monitoring and management. The architecture addresses a key gap in the field by showing how fragmented components can be organised into a coherent system-level approach.

The value of an integrative reference architecture is that IoT-based depression management cannot be achieved through sensing, prediction, or automation alone. Depression-related

data require context, interpretation, personalisation, safety mechanisms, human oversight, and responsible governance. *IoTMindCare* provides a foundation for bringing these requirements together within a single architecture.

The architecture-level contribution of *IoTMindCare* also reframes IoT-based depression management as a design problem that extends beyond technical functionality. The thesis shows that future system development needs to consider technical design alongside clinical relevance, user engagement, privacy, interoperability, stakeholder roles, and governance. In this sense, *IoTMindCare* provides a structured pathway from isolated technical prototypes towards more integrated and responsibly designed future systems.

13.4 Final Concluding Remarks

This thesis has shown that the future development of IoT-based depression systems requires more than improved sensing, more data, or more accurate prediction models. It requires an architecture capable of connecting data collection with contextual interpretation, behavioural engagement, personalised analytics, governed escalation, human oversight, and responsible system design.

IoTMindCare provides an architecture-level foundation for connecting data collection, contextual interpretation, behavioural engagement, personalised analytics, governed escalation, human oversight, and responsible system design. It integrates technical, behavioural, contextual, safety, and governance considerations that can inform the future development of safer, more personalised IoT-based depression monitoring and management systems. Further implementation, piloting, and validation are still required, but the thesis establishes a clear architecture-level contribution to IoT, digital mental health, and system architecture research.

In conclusion, this thesis contributes a structured and expert-informed foundation for moving IoT-based depression research beyond fragmented monitoring tools and isolated prediction models. It supports a pathway towards future IoT-based mental health systems that can be designed with stronger attention to personalisation, contextual awareness, safety, human oversight, and careful evaluation.

References

[1] Global Burden of Disease Collaborative Network. Global Burden of Disease Study 2023 (GBD 2023) Results, 2024. Accessed: 13 August 2025.

[2] World Health Organization. Depression, 2025. Accessed: 2026-06-07.

[3] World Health Organization. *ICD-11 for Mortality and Morbidity Statistics*. World Health Organization, Geneva, Switzerland, 2022.

[4] World Health Organization. Depression, Accessed 2024. Accessed on May 1, 2024.

[5] Andrew Sommerlad, Louise Marston, Jonathan Huntley, Gill Livingston, Gemma Lewis, Andrew Steptoe, and Daisy Fancourt. Social relationships and depression during the COVID-19 lockdown: longitudinal analysis of the COVID-19 Social Study. *Psychological medicine*, 52(15):3381–3390, 2022.

[6] Ministry of Health (New Zealand). 2024–25 new zealand health survey annual data explorer. https://minhealthnz.shinyapps.io/nz-health-survey-2024-25-annual-data-explorer/_w_3a97b811d50b4cc4b2c7afc9add22ecf/#!/explore-indicators, 2025. Interactive online data explorer presenting results for the 2024–25 New Zealand Health Survey. Accessed: 2026-02-17.

[7] Jeanette Tamm, Keisuke Takano, Leah Just, Thomas Ehring, Tabea Rosenkranz, and Johannes Kopf-Beck. Ecological momentary assessment versus weekly questionnaire assessment of change in depression. *Depression and Anxiety*, 2024(1):9191823, 2024.

[8] Luis San, Belen Arranz, Carlota Romans, Berta García, Marta Coromina, Sonia Ortiz, Miriam Vilaplana, Víctor Soto, Ruth Villaescusa, and Joan Alvaros. Patient-reported outcome measures to assess mental and physical health status, functionality, and quality of life in patients with major depression or schizophrenia. *Journal of Patient-Reported Outcomes*, 9(1):4, 2025.

[9] World Health Organization. *Depression and Other Common Mental Disorders: Global Health Estimates*. World Health Organization, Geneva, Switzerland, 2017.

[10] Katharine A Smith, Charlotte Blease, Maria Faurholt-Jepsen, Joseph Firth, Tom Van Daele, Carmen Moreno, Per Carlbring, Ulrich W Ebner-Priemer, Nikolaos Koutsouleris, Heleen Riper, et al. Digital mental health: challenges and next steps. *BMJ mental health*, 26(1), 2023.

[11] Raymond R Bond, Maurice D Mulvenna, Courtney Potts, Siobhan O'Neill, Edel Ennis, and John Torous. Digital transformation of mental health services. *Npj Mental Health Research*, 2(1):13, 2023.

[12] John Torous and Laura Weiss Roberts. Needed innovation in digital health and smartphone applications for mental health: transparency and trust. *JAMA psychiatry*, 74(5):437–438, 2017.

[13] Imogen E Leaning, Nessa Ikani, Hannah S Savage, Alex Leow, Christian Beckmann, Henricus G Ruhé, and Andre F Marquand. From smartphone data to clinically relevant predictions: A systematic review of digital phenotyping methods in depression. *Neuroscience & Biobehavioral Reviews*, 158:105541, 2024.

[14] Mahsa Sheikh, Meha Qassem, and Panicos A Kyriacou. Wearable, environmental, and smartphone-based passive sensing for mental health monitoring. *Frontiers in digital health*, 3:662811, 2021.

[15] ShiYing Shen, Wenhao Qi, Jianwen Zeng, Sixie Li, Xin Liu, Xiaohong Zhu, Chaoqun Dong, Bin Wang, Yankai Shi, Jiani Yao, et al. Passive sensing for mental health monitoring using machine learning with wearables and smartphones: Scoping review. *Journal of medical Internet research*, 27:e77066, 2025.

[16] Ivan Moura, Ariel Teles, Davi Viana, Jean Marques, Luciano Coutinho, and Francisco Silva. Digital phenotyping of mental health using multimodal sensing of multiple situations of interest: a systematic literature review. *Journal of biomedical informatics*, 138:104278, 2023.

[17] Danielle Currey and John Torous. Increasing the value of digital phenotyping through reducing missingness: a retrospective review and analysis of prior studies. *BMJ Mental Health*, 26(1), 2023.

[18] Carrie KY Lau, Anthony Saad, Bettina Camara, Dia Rahman, and Blanca Bolea-Alamanac. Acceptability of digital mental health interventions for depression and anxiety: systematic review. *Journal of medical Internet research*, 26:e52609, 2024.

[19] Katharine A Smith, Thomas Ward, Sinéad Lambe, Edoardo G Ostinelli, Charlotte Blease, Thomas Gant, Stefan M Gold, Emily A Holmes, Ivana Paccoud, Anastasia Vinnikova, et al. Engagement and attrition in digital mental health: current challenges and potential solutions. *NPJ digital medicine*, 8(1):398, 2025.

[20] Sanaz Zamani, Roopak Sinha, Minh Nguyen, and Samaneh Madanian. Enhancing emotional well-being with iot data solutions for depression: A systematic review. *IEEE Journal of Biomedical and Health Informatics*, 2025.

[21] Jessica M Lipschitz, Chelsea K Pike, Timothy P Hogan, Susan A Murphy, and Katherine E Burdick. The engagement problem: a review of engagement with digital mental health interventions and recommendations for a path forward: Lipschitz et al. *Current treatment options in psychiatry*, 10(3):119–135, 2023.

[22] Alan R Hevner, Salvatore T March, Jinsoo Park, and Sudha Ram. Design science in information systems research1. *MIS quarterly*, 28(1):75–106, 2004.

[23] Roel Wieringa. *Design science methodology for information systems and software engineering*. Springer, 2014.

[24] Sanaz Zamani, Minh Nguyen, and Roopak Sinha. Integrating environmental data for mental health monitoring: A data-driven iot-based approach. *Applied Sciences*, 15(2):912, 2025.

[25] Sanaz Zamani, Adnan Rostami, Minh Nguyen, Roopak Sinha, and Samaneh Madanian. A gamified ai-driven system for depression monitoring and management. *Applied Sciences*, 15(13):7088, 2025.

[26] Sanaz Zamani, Roopak Sinha, Samaneh Madanian, and Minh Nguyen. Iotmindcare: An integrative reference architecture for safe and personalized iot-based depression management. *Sensors*, 25(22):6994, 2025.

[27] Rebeka Amin, Simon Schreynemackers, Hannah Oppenheimer, Milica Petrovic, Ulrich Hegerl, and Hanna Reich. Use of mobile sensing data for longitudinal monitoring and prediction of depression severity: Systematic review. *Journal of Medical Internet Research*, 27:e57418, 2025.

[28] Hyun Woo Jung, Ilju Lee, Ok Kim, Seungjin Lee, Sujin Lee, Un Sun Chung, Jae-Hyun Kim, Sehwan Kim, Jung Won Kim, Ah Lahm Shin, et al. Key features of digital phenotyping for monitoring mental disorders: systematic review. *Journal of Medical Internet Research*, 27(1):e77331, 2025.

[29] Yannik Terhorst, Eva-Maria Messner, Kennedy Opoku Asare, Christian Montag, Christopher Kannen, and Harald Baumeister. Investigating smartphone-based sensing features for depression severity prediction: observation study. *Journal of Medical Internet Research*, 27:e55308, 2025.

[30] Yuting Zhan, Hongwei Liu, and Yu Wang. Digital phenotyping of depression: A multimodal passive sensing approach to identifying novel behavioral and physiological markers of treatment response. *Journal of Psychiatric Research*, 2025.

[31] Daniel Emden, Maike Richter, Astrid Chevance, Ramona Leenings, Julian Herpertz, Lara Gutfleisch, Anna Fleuchaus, Rogério Blitz, Vincent L Holstein, Janik Goltermann, et al. Scalable depression monitoring with smartphone speech using a multimodal benchmark and topic analysis. *npj Digital Medicine*, 2026.

[32] Elizabeth W Lampe, Amanda C Collins, Ahhyun Lee, Nicholas Enbar-Salo, Tess Z Griffin, Arvind Pillai, Michael V Heinz, Damien Lekkas, Matthew D Nemesure, Daniel M Mackin, et al. Interindividual differences in digital phenotypes of major depressive disorder: A passive sensing study using smartphone and wearable sensor data. *Behaviour research and therapy*, page 104986, 2026.

[33] Nur Hani Zainal and Peter F Hitchcock. Sleep disturbance recorded via wearable sensors predicts depression severity 9 years later. *Journal of Affective Disorders*, page 120426, 2025.

[34] Luyao Wang, Chenhan Wang, Chenyang Li, Toshiya Murai, Yicai Bai, Ziyan Song, Shuoyan Zhang, Qi Zhang, Yu Huang, Xiaoying Bi, et al. Ai-assisted multi-modal information for the screening of depression: a systematic review and meta-analysis. *NPJ Digital Medicine*, 8(1):523, 2025.

[35] Jiawen Liu, Junhui Wang, Zhaobin Wu, and Mohamad Ibrani Shahrimin Bin Adam Assim. The performance of wearable device–based artificial intelligence in detecting depression: Systematic review and meta-analysis. *JMIR Mental Health*, 13:e85319, 2026.

[36] Bikash Pradhan, Saugat Bhattacharyya, and Kunal Pal. Iot-based applications in healthcare devices. *Journal of healthcare engineering*, 2021(1):6632599, 2021.

[37] K Woodward, E Kanjo, D Brown, et al. Beyond mobile Apps: A survey of Technologies for mental well-being. 2020, 1905.

[38] Abinaya Gopalakrishnan, Raj Gururajan, Xujuan Zhou, Revathi Venkataraman, Ka Ching Chan, and Niall Higgins. A survey of autonomous monitoring systems in mental health. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, page e1527, 2024.

[39] Alieh Hajizadeh Saffar, Tiffany Katharine Mann, and Bahadorreza Ofoghi. Textual emotion detection in health: Advances and applications. *Journal of Biomedical Informatics*, 137:104258, 2023.

[40] Yekta Said Can, Bhargavi Mahesh, and Elisabeth André. Approaches, applications, and challenges in physiological emotion recognition—a tutorial overview. *Proceedings of the IEEE*, 2023.

[41] Ahmed Shihab Albahri, Jwan K Alwan, Zahraa K Taha, Sura F Ismail, Rula A Hamid, AA Zaidan, Osamah Shihab Albahri, BB Zaidan, Abdullah Hussein Alamoodi, and MA Alsalem. IoT-based telemedicine for disease prevention and health promotion: State-of-the-Art. *Journal of Network and Computer Applications*, 173:102873, 2021.

[42] Ahona Ghosh and Sharmistha Dey. "Sensing the Mind": An Exploratory Study About Sensors Used in E-Health and M-Health Applications for Diagnosis of Mental Health Condition. In *Efficient Data Handling for Massive Internet of Medical Things: Healthcare Data Analytics*, pages 269–292. Springer, 2021.

[43] Divya Gupta, MPS Bhatia, and Akshi Kumar. Real-time mental health analytics using iomt and social media datasets: research and challenges. In *Proceedings of the International Conference on Innovative Computing & Communication (ICICC)*, 2021.

[44] Tine Kolenik. Methods in digital mental health: smartphone-based assessment and intervention for stress, anxiety, and depression. In *Integrating Artificial Intelligence and IoT for Advanced Health Informatics: AI in the Healthcare Sector*, pages 105–128. Springer, 2022.

[45] Uichin Lee, Kyungsik Han, Hyunsung Cho, Kyong-Mee Chung, Hwajung Hong, Sung-Ju Lee, Youngtae Noh, Sooyoung Park, and John M Carroll. Intelligent positive computing with mobile, wearable, and IoT devices: Literature review and research directions. *Ad Hoc Networks*, 83:8–24, 2019.

[46] Sunitha Lingam. IoT Healthcare Applications. *The Fusion of Internet of Things, Artificial Intelligence, and Cloud Computing in Health Care*, pages 135–154, 2021.

[47] Monika Pahuja and Dinesh Kumar. Taxonomy of Various Routing Protocols of IoT-Based on Wireless Sensor Networks for Healthcare. *International Journal of Computer Networks and Applications*, pages 569–602, 2023.

[48] Thanveer Shaik, Xiaohui Tao, Lin Li, Haoran Xie, and Juan D Velásquez. A survey of multimodal information fusion for smart healthcare: Mapping the journey from data to wisdom. *Information Fusion*, page 102040, 2023.

[49] Barbara Kitchenham, O Pearl Brereton, David Budgen, Mark Turner, John Bailey, and Stephen Linkman. Systematic literature reviews in software engineering–a systematic literature review. *Information and software technology*, 51(1):7–15, 2009.

[50] Matthew J Page, Joanne E McKenzie, Patrick M Bossuyt, Isabelle Boutron, Tammy C Hoffmann, Cynthia D Mulrow, Larissa Shamseer, Jennifer M Tetzlaff, Elie A Akl, Sue E Brennan, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *International journal of surgery*, 88:105906, 2021.

[51] Adam M Chekroud, Julia Bondar, Jaime Delgadillo, Gavin Doherty, Akash Wasil, Marjolein Fokkema, Zachary Cohen, Danielle Belgrave, Robert DeRubeis, Raquel Iniesta, et al. The promise of machine learning in predicting treatment outcomes in psychiatry. *World Psychiatry*, 20(2):154–170, 2021.

[52] Enrique Garcia-Ceja, Michael Riegler, Tine Nordgreen, Petter Jakobsen, Ketil J Oedegaard, and Jim Tørresen. Mental health monitoring with multimodal sensing and machine learning: A survey. *Pervasive and Mobile Computing*, 51:1–26, 2018.

[53] Tony Robinson, Joan Condell, Elaine Ramsey, and Gerard Leavey. Self-management of subclinical common mental health disorders (anxiety, depression and sleep disorders) using wearable devices. *International journal of environmental research and public health*, 20(3):2636, 2023.

[54] Eduardo Teixeira, Hélder Fonseca, Florêncio Diniz-Sousa, Lucas Veras, Giorjines Boppre, José Oliveira, Diogo Pinto, Alberto Jorge Alves, Ana Barbosa, Romeu Mendes, et al. Wearable devices for physical activity and healthcare monitoring in elderly people: A critical review. *Geriatrics*, 6(2):38, 2021.

[55] Ayan Kumar, Abdul Quadir, J Christy Jackson, and Celestine Iwendi. Predicting and curing depression using long short term memory and global vector. *Computers, Materials & Continua*, 74(3), 2023.

[56] Garima Gupta, Sanskruti Joshi, and Sukanya Kulkarni. IoT based Depression Detection. In *2022 IEEE 3rd Global Conference for Advancement in Technology (GCAT)*, pages 1–7. IEEE, 2022.

[57] Lizeth-Guadalupe Machado-Jaimes, Martin Rogelio Bustamante-Bello, Amadeo-José Argüelles-Cruz, and Mariel Alfaro-Ponce. Development of an intelligent system for the monitoring and diagnosis of the well-being. *Sensors*, 22(24):9719, 2022.

[58] Prabha Sundaravadivel, Vividha Goyal, and Lakshman Tamil. I-Rise: An IoT-based semi-immersive affective monitoring framework for anxiety disorders. In *2020 IEEE International Conference on Consumer Electronics (ICCE)*, pages 1–5. IEEE, 2020.

[59] Farman Ali, Shaker El-Sappagh, SM Riazul Islam, Amjad Ali, Muhammad Attique, Muhammad Imran, and Kyung-Sup Kwak. An intelligent healthcare monitoring framework using wearable sensors and social networking data. *Future Generation Computer Systems*, 114:23–43, 2021.

[60] Sihao Yang, Bin Gao, Long Jiang, Jikun Jin, Zhao Gao, Xiaole Ma, and Wai Lok Woo. IoT structured long-term wearable social sensing for mental wellbeing. *IEEE Internet of Things Journal*, 6(2):3652–3662, 2018.

[61] Prabha Sundaravadivel, Paola Salvatore, and Premananda Indic. M-SID: an IoT-based edge-intelligent framework for suicidal ideation detection. In *2020 IEEE 6th World Forum on Internet of Things (WF-IoT)*, pages 1–6. IEEE, 2020.

[62] Asal Yunusova, Jocelyn Lai, Alexander P Rivera, Sirui Hu, Sina Labbaf, Amir M Rahmani, Nikil Dutt, Ramesh C Jain, Jessica L Borelli, et al. Assessing the mental health of emerging adults through a mental health app: Protocol for a prospective pilot study. *JMIR Research Protocols*, 10(3):e25775, 2021.

[63] Akshi Kumar, Kapil Sharma, and Aditi Sharma. Hierarchical deep neural network for mental stress state detection using IoT based biomarkers. *Pattern Recognition Letters*, 145:81–87, 2021.

[64] Timon Elmer and Christoph Stadtfeld. Depressive symptoms are associated with social isolation in face-to-face interaction networks. *Scientific reports*, 10(1):1444, 2020.

[65] Jian Chen, Yuzhu Hu, Qifeng Lai, Wei Wang, Junxin Chen, Han Liu, Gautam Srivastava, Ali Kashif Bashir, and Xiping Hu. IIFDD: Intra and inter-modal fusion for depression detection with multi-modal information from Internet of Medical Things. *Information Fusion*, 102:102017, 2024.

[66] Leandro Y Mano, Bruno S Faiçal, Luis HV Nakamura, Pedro H Gomes, Giampaolo L Libralon, Rodolfo I Meneguete, P Rocha Geraldo Filho, Gabriel T Giancristofaro, Gustavo Pessin, Bhaskar Krishnamachari, et al. Exploiting IoT technologies for enhancing Health Smart Homes through patient identification and emotion recognition. *Computer Communications*, 89:178–190, 2016.

[67] Jiajue Ou, Huiguang Liang, and Hwee Xian Tan. Identifying elderlies at risk of becoming more depressed with internet-of-things. In *Human Aspects of IT for the Aged Population. Applications in Health, Assistance, and Entertainment: 4th International Conference, ITAP 2018, Held as Part of HCI International 2018, Las Vegas, NV, USA, July 15–20, 2018, Proceedings, Part II 4*, pages 348–361. Springer, 2018.

[68] Ivan Rodrigues de Moura, Ariel Soares Teles, Markus Endler, Luciano Reis Coutinho, and Francisco José da Silva E Silva. Recognizing context-aware human sociability patterns using pervasive monitoring for supporting mental health professionals. *Sensors*, 21(1):86, 2020.

[69] Ali Mahdi, Ahmed Kazim, Alia AlHammadi, and Vinod Pangracious. EmoGo: A Smart Wearable IoT System for Human Emotion Detection. In *2022 14th Annual Undergraduate Research Conference on Applied Computing (URC)*, pages 1–4. IEEE, 2022.

[70] Joel Mathew Jose, Jovel Varghese Jose, and Chaitanya Vijaykumar Mahamuni. Multi-biosensor based wireless body area networks (wban) for critical health monitoring of patients in mental health care centers: An interdisciplinary study. *International Journal of Research in Engineering, Science and Management*, 3, 2020.

[71] Lázaro Hardy, Diego De Castro, Guilherme Agnolin, Gabriel Ronney da Silva, and Marcela Purificação. A New Framework for Gaining Emotional Health Knowledge Through Virtual Reality-IoT Technology. In *European Conference on Knowledge Management*, volume 24, pages 518–524, 2023.

[72] Arthur Gatouillat, Youakim Badr, Bertrand Massot, and Ervin Sejdić. Internet of medical things: A review of recent contributions dealing with cyber-physical systems in medicine. *IEEE internet of things journal*, 5(5):3810–3822, 2018.

[73] Ana Carolina Borges Monteiro, Reinaldo Padilha França, Rangel Arthur, and Yuzo Iano. An overview of medical internet of things, artificial intelligence, and cloud computing employed in health care from a modern panorama. *The Fusion of Internet of Things, Artificial Intelligence, and Cloud Computing in Health Care*, pages 3–23, 2021.

[74] Mijeong Kang and Kyunghwan Chai. Wearable sensing systems for monitoring mental health. *Sensors*, 22(3):994, 2022.

[75] Jamie T Mullins and Corey White. Temperature and mental health: Evidence from the spectrum of mental health outcomes. *Journal of health economics*, 68:102240, 2019.

[76] Anand Singh Rajawat, Romil Rawat, Kanishk Barhanpurkar, Rabindra Nath Shaw, and Ankush Ghosh. Depression detection for elderly people using ai robotic systems leveraging the nelder–mead method. In *Artificial Intelligence for Future Generation Robotics*, pages 55–70. Elsevier, 2021.

[77] M Srividya, Subramaniam Mohanavalli, and Natarajan Bhalaji. Behavioral modeling for mental health using machine learning algorithms. *Journal of medical systems*, 42:1–12, 2018.

[78] Yaoxi Shi, Peter Koval, Vassilis Kostakos, Jorge Goncalves, and Greg Wadley. "instant happiness": Smartphones as tools for everyday emotion regulation. *International Journal of Human-Computer Studies*, 170:102958, 2023.

[79] Xinyu Shui, Mi Zhang, Zhuoran Li, Xin Hu, Fei Wang, and Dan Zhang. A dataset of daily ambulatory psychological and physiological recording for emotion research. *Scientific data*, 8(1):161, 2021.

[80] Juan Abdon Miranda-Correa, Mojtaba Khomami Abadi, Nicu Sebe, and Ioannis Patras. Amigos: A dataset for affect, personality and mood research on individuals and groups. *IEEE transactions on affective computing*, 12(2):479–493, 2018.

[81] Karan Sharma, Claudio Castellini, Egon L van den Broek, Alin Albu-Schaeffer, and Friedhelm Schwenker. A dataset of continuous affect annotations and physiological signals for emotion analysis. *Scientific data*, 6(1):196, 2019.

[82] Mário WL Moreira, Joel JPC Rodrigues, Neeraj Kumar, Kashif Saleem, and Igor V Illin. Postpartum depression prediction through pregnancy data analysis for emotion-aware smart systems. *Information Fusion*, 47:23–31, 2019.

[83] Sheona Sinha, Shubham Kr Mishra, and Saurabh Bilgaiyan. Emotion analysis to provide counseling to students fighting from depression and anxiety by using cctv surveillance. In *Machine Learning and Information Processing: Proceedings of ICMLIP 2019*, pages 81–94. Springer, 2020.

[84] Surendrabikram Thapa, Awishkar Ghimire, Surabhi Adhikari, Akash Kumar Bhoi, and Paolo Barsocchi. Cognitive Internet of Things (IoT) and computational intelligence for mental well-being. In *Cognitive and Soft Computing Techniques for the Analysis of Healthcare Data*, pages 59–77. Elsevier, 2022.

[85] Akshi Kumar, Kapil Sharma, and Aditi Sharma. Memor: A multimodal emotion recognition using affective biomarkers for smart prediction of emotional health for people analytics in smart industries. *Image and Vision Computing*, 123:104483, 2022.

[86] Anand Singh Rajawat, Kanishk Barhanpurkar, Rabindra Nath Shaw, and Ankush Ghosh. Risk detection in wireless body sensor networks for health monitoring using hybrid deep learning. In *Innovations in Electrical and Electronic Engineering: Proceedings of ICEEE 2021*, pages 683–696. Springer, 2021.

[87] Amal Tawakuli and Thomas Engel. Towards normalizing the design phase of data preprocessing pipelines for iot data. In *2022 IEEE International Conference on Big Data (Big Data)*, pages 4589–4594. IEEE, 2022.

[88] Amelie Gyrard, Kasia Tabeau, Laura Fiorini, Antonio Kung, Eloise Senges, Marleen De Mul, Francesco Giuliani, Delphine Lefebvre, Hiroshi Hoshino, Isabelle Fabbricotti, et al. Knowledge Engineering Framework for IoT Robotics Applied to Smart Healthcare and Emotional Well-Being. *International Journal of Social Robotics*, pages 1–28, 2021.

[89] Cristina Mele, Marialuisa Marzullo, Irene Di Bernardo, Tiziana Russo-Spena, Roberta Massi, Alessandra La Salandra, and Stefania Cialabrini. A smart tech lever to augment caregivers' touch and foster vulnerable patient engagement and well-being. *Journal of Service Theory and Practice*, 32(1):52–74, 2022.

[90] Amandeep Kaur and Deepak Sharma. Design of decision support system for diagnosis of psychological disorder. *International Journal of Computer Applications and Information Technology*, 12(2):352–363, 2020.

[91] Reham Alhejaili and Akram Alomainy. The use of wearable technology in providing assistive solutions for mental well-being. *Sensors*, 23(17):7378, 2023.

[92] Robert H Dolin, Liora Alschuler, Sandy Boyer, Calvin Beebe, Fred M Behlen, Paul V Biron, and Amnon Shabo. HL7 clinical document architecture, release 2. *Journal of the American Medical Informatics Association*, 13(1):30–39, 2006.

[93] Manuella Lech Cantuaria, Jørgen Brandt, and Victoria Blanes-Vidal. Exposure to multiple environmental stressors, emotional and physical well-being, and self-rated health: An analysis of relationships using latent variable structural equation modelling. *Environmental Research*, 227:115770, 2023.

[94] Georgianne Tiu Hawkins, Chloe S Chung, Marci F Hertz, and Noehealani Antolin. The school environment and physical and social-emotional well-being: implications for students and school employees. *Journal of School Health*, 93(9):799–812, 2023.

[95] Paolo Cianconi, Sophia Betrò, and Luigi Janiri. The impact of climate change on mental health: a systematic descriptive review. *Frontiers in psychiatry*, 11:490206, 2020.

[96] Marvin Bundo, Martin Preisig, Kathleen Merikangas, Jennifer Glaus, Julien Vaucher, Gérard Waeber, Pedro Marques-Vidal, Marie-Pierre F Strippoli, Thomas Müller, Oscar Franco, et al. How ambient temperature affects mood: an ecological momentary assessment study in switzerland. *Environmental Health*, 22(1):52, 2023.

[97] Mir Sajjad Hussain Talpur, Abdul Ahad Abro, Mansoor Ebrahim, Irfan Ali Kandhro, Selvakumar Manickam, Shams Ul Arfeen Laghari, Abdulhalim Dandoush, and Mueen Uddin. Illuminating healthcare management: A comprehensive review of iot-enabled chronic disease monitoring. *IEEE Access*, 12:48189–48209, 2024.

[98] Fatma M Talaat and Rana Mohamed El-Balka. Stress monitoring using wearable sensors: Iot techniques in medical field. *Neural Computing and Applications*, 35(25):18571–18584, 2023.

[99] Fernando Florido Ngu, Ilan Kelman, Jonathan Chambers, and Sonja Ayeb-Karlsson. correlating heatwaves and relative humidity with suicide (fatal intentional self-harm). *Scientific reports*, 11(1):22175, 2021.

[100] American Psychiatric Association. Extreme heat can take a toll on mental health, 2023. Accessed: 18 December 2024.

[101] Moustaq Karim Khan Rony and Hasnat M Alamgir. High temperatures on mental health: Recognizing the association and the need for proactive strategies—a perspective. *Health science reports*, 6(12):e1729, 2023.

[102] Clemens Noelke, Mark McGovern, Daniel J Corsi, Marcia P Jimenez, Ari Stern, Ian Sue Wing, and Lisa Berkman. Increasing ambient temperature reduces emotional well-being. *Environmental research*, 151:124–129, 2016.

[103] Public Health Communication Centre. Unaffordable home heating increases risk of severe mental distress, 2023. Accessed: 20 December 2024.

[104] World Health Organization. Mental disorders, 2022. Accessed: 18 December 2024.

[105] Sebastian Bauhoff. Self-report bias in estimating cross-sectional and treatment effects. In *Encyclopedia of quality of life and well-being research*, pages 6277–6279. Springer, 2024.

[106] Emese Sükei, Agnes Norbury, M Mercedes Perez-Rodriguez, Pablo M Olmos, and Antonio Artés. Predicting emotional states using behavioral markers derived from passively sensed data: data-driven machine learning approach. *JMIR mHealth and uHealth*, 9(3):e24465, 2021.

[107] Marlies Houben, Wim Van Den Noortgate, and Peter Kuppens. The relation between short-term emotion dynamics and psychological well-being: A meta-analysis. *Psychological bulletin*, 141(4):901, 2015.

[108] G Fretes, C Llurba, and R Palau. Exploring classroom temperature and humidity on students' emotions through iot and image processing. *Journal of Infrastructure, Policy and Development*, 8(7):1–14, 2024.

[109] Mehdi Dadkhah, Mohammad Mehraeen, Fariborz Rahimnia, and Khalil Kimiafar. Use of internet of things for chronic disease management: an overview. *Journal of Medical Signals & Sensors*, 11(2):138–157, 2021.

[110] World Health Organization. Urban green spaces and health. Technical Report WHO/EURO: 2016-3352-43111-60341, World Health Organization, Regional Office for Europe, Copenhagen, Denmark, 2016.

[111] Theresa HM Moore, Joanna M Kesten, Jose A López-López, S Ijaz, A McAleenan, Alison Richards, Selena Gray, Jelena Savović, and Suzanne Audrey. The effects of changes to the built environment on the mental health and well-being of adults: Systematic review. *Health & place*, 53:237–257, 2018.

[112] Xin-qi Liu, Jia Huang, Chao Song, Tian-liang Zhang, Yong-ping Liu, and Li Yu. Neurodevelopmental toxicity induced by pm2. 5 exposure and its possible role in neurodegenerative and mental disorders. *Human & Experimental Toxicology*, 42:09603271231191436, 2023.

[113] Yanhong Peng, Xiaoyan Yang, Dongze Li, Zisu Ma, Zhuochao Liu, Xuehang Bai, and Zebing Mao. Predicting flow status of a flexible rectifier using cognitive computing. *Expert Systems with Applications*, 264:125878, 2025.

[114] Yanhong Peng, Yuxin Wang, Fangchao Hu, Miao He, Zebing Mao, Xia Huang, and Jun Ding. Predictive modeling of flexible ehd pumps using kolmogorov–arnold networks. *Biomimetic Intelligence and Robotics*, 4(4):100184, 2024.

[115] Sandeep Reddy, John Fox, and Maulik P Purohit. Artificial intelligence-enabled healthcare delivery. *Journal of the Royal Society of Medicine*, 112(1):22–28, 2019.

[116] Andre Esteva, Alexandre Robicquet, Bharath Ramsundar, Volodymyr Kuleshov, Mark DePristo, Katherine Chou, Claire Cui, Greg Corrado, Sebastian Thrun, and Jeff Dean. A guide to deep learning in healthcare. *Nature medicine*, 25(1):24–29, 2019.

[117] Senerath Mudalige Don Alexis Chinthaka Jayatilake and Gamage Upeksha Ganegoda. Involvement of machine learning tools in healthcare decision making. *Journal of healthcare engineering*, 2021(1):6679512, 2021.

[118] Steven A Gray, Stefan Gray, Linda J Cox, and Sarah Henly-Shepard. Mental modeler: a fuzzy-logic cognitive mapping modeling tool for adaptive environmental management. In *2013 46th Hawaii international conference on system sciences*, pages 965–973. IEEE, 2013.

[119] Ch Venkateswara Rao, MKV Subba Reddy, J Kishore, RV Phani Sirisha, B Srinivasa Seshagiri Rao, and Himakiran Killamsetti. Optimizing indoor air quality: a fuzzy logic-based control system. In *2023 2nd International Conference on Automation, Computing and Renewable Systems (ICACRS)*, pages 1053–1057. IEEE, 2023.

[120] Jin Pan, Julian Tang, Miguela Caniza, Jean-michel Heraud, Evelyn Koay, Hong Kai Lee, Chun Kiat Lee, Yuguo Li, Alejandra Nava Ruiz, Carlos Francisco Santillan-Salas, et al. Correlating indoor and outdoor temperature and humidity in a sample of buildings in tropical climates. *Indoor air*, 31(6):2281–2295, 2021.

[121] Juan Botero-Valencia, Luis Castano-Londono, and David Marquez-Viloria. Indoor temperature and relative humidity dataset of controlled and uncontrolled environments. *Data*, 7(6):81, 2022.

[122] Jian-Lei Kong, Xiao-Meng Fan, Xue-Bo Jin, Ting-Li Su, Yu-Ting Bai, Hui-Jun Ma, and Min Zuo. Bmae-net: A data-driven weather prediction network for smart agriculture. *Agronomy*, 13(3):625, 2023.

[123] Judith Borghouts, Elizabeth Eikey, Gloria Mark, Cinthia De Leon, Stephen M Schueller, Margaret Schneider, Nicole Stadnick, Kai Zheng, Dana Mukamel, and Dara H Sorkin. Barriers to and facilitators of user engagement with digital mental health interventions: systematic review. *Journal of medical Internet research*, 23(3):e24387, 2021.

[124] Raymond W Lam, Sidney H Kennedy, Roger S McIntyre, and Atul Khullar. Cognitive dysfunction in major depressive disorder: effects on psychosocial functioning and implications for treatment. *The Canadian Journal of Psychiatry*, 59(12):649–654, 2014.

[125] John Torous, Jessica Lipschitz, Michelle Ng, and Joseph Firth. Dropout rates in clinical trials of smartphone apps for depressive symptoms: a systematic review and meta-analysis. *Journal of affective disorders*, 263:413–419, 2020.

[126] Amit Baumel, Frederick Muench, Stav Edan, and John M Kane. Objective user engagement with mental health apps: systematic search and panel-based usage analysis. *Journal of medical Internet research*, 21(9):e14567, 2019.

[127] Inbal Nahum-Shani, Shawna N Smith, Bonnie J Spring, Linda M Collins, Katie Witkiewitz, Ambuj Tewari, and Susan A Murphy. Just-in-time adaptive interventions (jitais) in mobile health: key components and design principles for ongoing health behavior support. *Annals of behavioral medicine*, pages 1–17, 2016.

[128] Daniel ZQ Gan, Lauren McGillivray, Mark E Larsen, Helen Christensen, and Michelle Torok. Technology-supported strategies for promoting user engagement with digital mental health interventions: a systematic review. *Digital health*, 8:20552076221098268, 2022.

[129] Theresa M Fleming, Lynda Bavin, Karolina Stasiak, Eve Hermansson-Webb, Sally N Merry, Colleen Cheek, Mathijs Lucassen, Ho Ming Lau, Britta Pollmuller, and Sarah Hetrick. Serious games and gamification for mental health: current status and promising directions. *Frontiers in psychiatry*, 7:215, 2017.

[130] Saul Shiffman, Arthur A Stone, and Michael R Hufford. Ecological momentary assessment. *Annual Review of Clinical Psychology*, 4:1–32, 2008.

[131] Silvan Hornstein, Kirsten Zantvoort, Ulrike Lueken, Burkhardt Funk, and Kevin Hilbert. Personalization strategies in digital mental health interventions: a systematic review and conceptual framework for depressive symptoms. *Frontiers in digital health*, 5:1170002, 2023.

[132] Mark Erik Larsen, Kit Huckvale, Jennifer Nicholas, John Torous, Louise Birrell, Emily Li, and Bill Reda. Using science to sell apps: evaluation of mental health app store quality claims. *NPJ digital medicine*, 2(1):18, 2019.

[133] Mei Yi Ng, Jennifer A Frederick, Aaron J Fisher, Nicholas B Allen, Jeremy W Pettit, and Dana L McMakin. Identifying person-specific drivers of depression in adolescents: Protocol for a smartphone-based ecological momentary assessment and passive sensing study. *JMIR Research Protocols*, 13(1):e43931, 2024.

[134] Ann Marie Roepke, Sara R Jaffee, Olivia M Riffle, Jane McGonigal, Rose Broome, and Bez Maxwell. Randomized controlled trial of superbetter, a smartphone-based/internet-based self-help tool to reduce depressive symptoms. *Games for health journal*, 4(3):235–246, 2015.

[135] Sally N Merry, Karolina Stasiak, Matthew Shepherd, Chris Frampton, Theresa Fleming, and Mathijs FG Lucassen. The effectiveness of sparx, a computerised self help intervention for adolescents seeking help for depression: randomised controlled non-inferiority trial. *Bmj*, 344, 2012.

[136] Akash R Wasil, Sarah Gillespie, Raveena Patel, Annemarie Petre, Katherine E Venturo-Conerly, Rebecca M Shingleton, John R Weisz, and Robert J DeRubeis. Reassessing evidence-based content in popular smartphone apps for depression and anxiety: Developing and applying user-adjusted analyses. *Journal of Consulting and Clinical Psychology*, 88(11):983, 2020.

[137] Gustavo F Tondello and Lennart E Nacke. Player characteristics and video game preferences. In *Proceedings of the annual symposium on computer-human interaction in play*, pages 365–378, 2019.

[138] Philip Cheng, Annemarie I Luik, Cynthia Fellman-Couture, Edward Peterson, Christine LM Joseph, Gabriel Tallent, Kieulinh Michelle Tran, Brian K Ahmedani, Timothy Roehrs, Thomas Roth, et al. Efficacy of digital cbt for insomnia to reduce depression across demographic groups: a randomized trial. *Psychological medicine*, 49(3):491–500, 2019.

[139] John Torous, Keris Jän Myrick, Natali Rauseo-Ricupero, Joseph Firth, et al. Digital mental health and covid-19: using technology today to accelerate the curve on access and quality tomorrow. *JMIR mental health*, 7(3):e18848, 2020.

[140] Adam DI Kramer, Jamie E Guillory, and Jeffrey T Hancock. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790, 2014.

[141] Fabian Wahle, Tobias Kowatsch, Elgar Fleisch, Michael Rufer, Steffi Weidt, et al. Mobile sensing and support for people with depression: a pilot trial in the wild. *JMIR mHealth and uHealth*, 4(3):e5960, 2016.

[142] Michael Rosenblum, Peter Miller, Benjamin Reist, Elizabeth A Stuart, Michael Thieme, and Thomas A Louis. Adaptive design in surveys and clinical trials: similarities, differences and opportunities for cross-fertilization. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 182(3):963–982, 2019.

[143] Aki Rintala, Martien Wampers, Inez Myin-Germeys, and Wolfgang Viechtbauer. Response compliance and predictors thereof in studies using the experience sampling method. *Psychological assessment*, 31(2):226, 2019.

[144] David C Mohr, Mi Zhang, and Stephen M Schueller. Personal sensing: understanding mental health using ubiquitous sensors and machine learning. *Annual review of clinical psychology*, 13(1):23–47, 2017.

[145] Kathleen Kara Fitzpatrick, Alison Darcy, and Molly Vierhile. Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (woebot): a randomized controlled trial. *JMIR mental health*, 4(2):e7785, 2017.

[146] Anant Kumar and K Rajasekharan Nayar. Covid 19 and its mental health consequences, 2021.

[147] Andreas Holzinger, Chris Biemann, Constantinos S Pattichis, and Douglas B Kell. What do we need to build explainable ai systems for the medical domain? *arXiv preprint arXiv:1712.09923*, 2017.

[148] Zhenguang Zhong and Zhixuan Wang. Intelligent depression prevention via llm-based dialogue analysis: Overcoming the limitations of scale-dependent diagnosis through precise emotional pattern recognition. *arXiv preprint arXiv:2504.16504*, 2025.

[149] Jiahao Qiu, Yinghui He, Xinzhe Juan, Yimin Wang, Yuhan Liu, Zixin Yao, Yue Wu, Xun Jiang, Ling Yang, and Mengdi Wang. Emoagent: Assessing and safeguarding human-ai interaction for mental health safety. *arXiv preprint arXiv:2504.09689*, 2025.

[150] Yuchen Hou, Wencai Chen, Sijing Chen, Xiujun Liu, Yingxin Zhu, Xiuqing Cui, and Bin Cao. Associations between indoor thermal environment assessment, mental health, and insomnia in winter. *Sustainable Cities and Society*, 114:105751, 2024.

[151] Nick Obradovich, Robyn Migliorini, Martin P Paulus, and Iyad Rahwan. Empirical evidence of mental health risks posed by climate change. *Proceedings of the National Academy of Sciences*, 115(43):10953–10958, 2018.

[152] Sherri Melrose. Seasonal affective disorder: an overview of assessment and treatment approaches. *Depression research and treatment*, 2015(1):178564, 2015.

[153] Katherine McLeod, Petros Spachos, and Konstantinos N Plataniotis. Smartphone-based wellness assessment using mobile environmental sensors. *IEEE Systems Journal*, 15(2):1989–1999, 2020.

[154] Laura SP Bloomfield, Mikaela I Fudolig, Julia Kim, Jordan Llorin, Juniper L Lovato, Ellen W McGinnis, Ryan S McGinnis, Matt Price, Taylor H Ricketts, Peter Sheridan Dodds, et al. Predicting stress in first-year college students using sleep data from wearable devices. *PLOS Digital Health*, 3(4):e0000473, 2024.

[155] Chenyu Tang, Wentian Yi, Muzi Xu, Yuxuan Jin, Zibo Zhang, Xuhang Chen, Caizhi Liao, Mengtian Kang, Shuo Gao, Peter Smielewski, et al. A deep learning–enabled smart garment for accurate and versatile monitoring of sleep conditions in daily life. *Proceedings of the National Academy of Sciences*, 122(7):e2420498122, 2025.

[156] Lakshan N Fonseka and Benjamin KP Woo. Wearables in schizophrenia: update on current and future clinical applications. *JMIR mHealth and uHealth*, 10(4):e35600, 2022.

[157] Leon Voon. How human-centered design can help public agencies design better digital services, August 2022. Accessed: 2025-05-16.

[158] Courtney Garvin. Transforming student wellbeing support with amazon bedrock and sxp.ai, May 2025. Accessed: 2025-05-16.

[159] Daniel S Pine and John A Cohen. Cognitive behavioral therapy, serotonin, and neural mechanisms of depression: the role of learning. *Psychological Medicine*, 41(7):1219–1230, 2011.

[160] Bradley E Holtz, Jennifer A Murray, and Todd Hershey. Serious games for cognitive training in mental health conditions: a systematic review. *JMIR Mental Health*, 8(2):e22007, 2021.

[161] Dominika Ochnik, Bartłomiej Buława, Paulina Nagel, Marek Gachowski, and Marcin Budziński. Urbanization, loneliness and mental health model-a cross-sectional network analysis with a representative sample. *Scientific Reports*, 14(1):24974, 2024.

[162] Ying Zhang, Xiaocan Jia, Yongli Yang, Na Sun, Shuyan Shi, and Wei Wang. Change in the global burden of depression from 1990-2019 and its prediction for 2030. *Journal of psychiatric research*, 178:16–22, 2024.

[163] Gregory E Simon, Nathalie Moise, and David C Mohr. Management of depression in adults: a review. *Jama*, 332(2):141–152, 2024.

[164] Qiong Wu, Xu Chen, Zhi Zhou, and Junshan Zhang. Fedhome: Cloud-edge based personalized federated learning for in-home health monitoring. *IEEE Transactions on Mobile Computing*, 21(8):2818–2832, 2020.

[165] Deepti Gupta, Olumide Kayode, Smriti Bhatt, Maanak Gupta, and Ali Saman Tosun. Hierarchical federated learning based anomaly detection using digital twins for smart healthcare. In *2021 IEEE 7th international conference on collaboration and internet computing (CIC)*, pages 16–25. IEEE, 2021.

[166] Humberto Cervantes and Rick Kazman. *Designing software architectures: a practical approach*. Addison-Wesley Professional, 2024.

[167] SA Angelov, Jos JM Trienekens, and PWPJ Grefen. Extending and adapting the architecture tradeoff analysis method for the evaluation of software reference architectures. 2014.

[168] SMART BEAR Consortium. Smart big data platform to offer evidence-based personalised support for healthy and independent living at home. https://smart-bear.eu/, 2020. Horizon 2020 Project. Accessed: 9 April 2026.

[169] Emma L Tonkin, Michael Holmes, Hao Song, Niall Twomey, Tom Diethe, Meelis Kull, Miquel Perello Nieto, Massimo Camplani, Sion Hannuna, Xenofon Fafoutis, et al. A multi-sensor dataset with annotated activities of daily living recorded in a residential setting. *Scientific Data*, 10(1):162, 2023.

[170] Ariza-Colpas Paola Patricia, Pacheco-Cuentas Rosberg, Shariq Butt-Aziz, Pineres-Melo Marlon Alberto, Morales-Ortega Roberto-Cesar, Urina-Triana Miguel, and Sumera Naz. Semi-supervised ensemble learning for human activity recognition in casas kyoto dataset. *Heliyon*, 10(8), 2024.

[171] Vinicius Facco Rodrigues, Rodrigo da Rosa Righi, Cristiano André da Costa, Felipe André Zeiser, Bjoern Eskofier, Andreas Maier, and Daeyoung Kim. Digital health in smart cities: Rethinking the remote health monitoring architecture on combining edge, fog, and cloud. *Health and Technology*, 13(3):449–472, 2023.

[172] David Benrimoh, Robert Fratila, Sonia Israel, and Kelly Perlman. Deep learning: A new horizon for personalized treatment of depression? *McGill Journal of Medicine*, 16(1), 2018.

[173] Javed Aalam, Syed Naseer Ahmad Shah, and Rafat Parveen. Personalized healthcare services for assisted living in healthcare 5.0. In *Ambient Assisted Living*, pages 203–222. CRC Press, 2025.

[174] Mladjan Jovanovic, Goran Mitrov, Eftim Zdravevski, Petre Lameski, Sara Colantonio, Martin Kampel, Hilda Tellioglu, and Francisco Florez-Revuelta. Ambient assisted living: scoping review of artificial intelligence models, domains, technology, and concerns. *Journal of Medical Internet Research*, 24(11):e36553, 2022.

[175] Yuezhou Zhang, Amos A Folarin, Callum Stewart, Heet Sankesara, Yatharth Ranjan, Pauline Conde, Akash Roy Choudhury, Shaoxiong Sun, Zulqarnain Rashid, and Richard JB Dobson. An explainable anomaly detection framework for monitoring depression and anxiety using consumer wearable devices. *arXiv preprint arXiv:2505.03039*, 2025.

[176] Anna Li, Eliane Bodanese, Stefan Poslad, Penghui Chen, Jun Wang, Yonglei Fan, and Tianwei Hou. A contactless health monitoring system for vital signs monitoring, human activity recognition, and tracking. *IEEE Internet of Things Journal*, 11(18):29275–29286, 2023.

[177] Health Level Seven International (HL7). HL7 FHIR (Fast Healthcare Interoperability Resources). Technical Report Release 5.0.0 (R5 – STU), HL7 International, March 2023. Accessed: 2025-06-29.

[178] openEHR International. openEHR – the future of digital health is open. https://openehr.org/, 2025. Accessed: 2025-06-29.

[179] FIWARE Foundation. FIWARE: Open apis for open minds — open-source smart applications platform. https://www.fiware.org/, 2025. Accessed: 2025-06-29.

[180] IEEE std 11073-10701–2022 / iso/ieee 11073-10701:2024: Health informatics—device interoperability—part 10701: Point-of-care medical device communication—metric provisioning by participants in a service-oriented device connectivity (sdc) system, September 2022. Published: 22 September 2023; ISO adopted: 2024-09.

[181] Judith J Prochaska, Erin A Vogel, Amy Chieng, Matthew Kendra, Michael Baiocchi, Sarah Pajarito, and Athena Robinson. A therapeutic relational agent for reducing problematic substance use (woebot): development and usability study. *Journal of medical Internet research*, 23(3):e24850, 2021.

[182] Touchkin eServices Pvt Ltd. Wysa: Your ai-powered mental wellness companion, 2025. Accessed: 2025-05-28.

[183] BetterHelp Editorial Team. How effective is online counseling for depression?, 2025. Accessed: 2025-04-29.

[184] Roble Ridge Software LLC. Moodfit: Tools & insights for your mental health, 2025. Accessed: 2025-04-29.

[185] Healthify NZ. T2 mood tracker app, 2025. Accessed: 2025-04-29.

[186] Software Engineering Design Research Group. Quality attribute scenario template – design practice repository, n.d. Accessed: 2025-05-18.

[187] Khayyam Akhtar, Muhammad Usman Yaseen, Muhammad Imran, Sohaib Bin Altaf Khattak, and Moustafa M Nasralla. Predicting inmate suicidal behavior with an interpretable ensemble machine learning approach in smart prisons. *PeerJ Computer Science*, 10:e2051, 2024.

[188] Stamatis Karnouskos, Roopak Sinha, Paulo Leitão, Luis Ribeiro, and Thomas I Strasser. The applicability of iso/iec 25023 measures to the integration of agents and automation systems. In *IECON 2018-44th Annual Conference of the IEEE Industrial Electronics Society*, pages 2927–2934. IEEE, 2018.

[189] Omar Cheikhrouhou, Khaleel Mershad, Faisal Jamil, Redowan Mahmud, Anis Koubaa, and Sanaz Rahimi Moosavi. A lightweight blockchain and fog-enabled secure remote patient monitoring system. *Internet of Things*, 22:100691, 2023.

[190] Philippe B Kruchten. The 4+ 1 view model of architecture. *IEEE software*, 12(6):42–50, 2002.

[191] Somayeh Nasiri, Farahnaz Sadoughi, Afsaneh Dehnad, Mohammad Hesam Tadayon, and Hossein Ahmadi. Layered architecture for internet of things-based healthcare system. *Informatica (Slovenia)*, 45(4), 2021.

[192] Francesco Lamonaca, Carmelo Scuro, Domenico Grimaldi, Renato Sante Olivito, Paolo Francesco Sciammarella, and Domenico Luca Carnì. A layered iot-based architecture for a distributed structural health monitoring system system. *Acta Imeko*, 8(2):45–52, 2019.

[193] Jinyoung Choi, Soomin Lee, Seonyoung Kim, Dongil Kim, and Hyungshin Kim. Depressed mood prediction of elderly people with a wearable band. *Sensors*, 22(11):4174, 2022.

[194] Leonard Richardson and Sam Ruby. *RESTful web services*. " O'Reilly Media, Inc.", 2008.

[195] Mohammad Aazam, Sherali Zeadally, and Khaled A Harras. Deploying fog computing in industrial internet of things and industry 4.0. *IEEE Transactions on Industrial Informatics*, 14(10):4674–4682, 2018.

[196] Mohammad Abdur Razzaque, Marija Milojevic-Jevric, Andrei Palade, and Siobhán Clarke. Middleware for internet of things: a survey. *IEEE Internet of things journal*, 3(1):70–95, 2015.

[197] CarePredict. Carepredict: Ai-powered senior care platform, 2025. Accessed: 2025-06-06.

[198] Ashkan Yousefpour, Caleb Fung, Tam Nguyen, Krishna Kadiyala, Fatemeh Jalali, Amirreza Niakanlahiji, Jian Kong, and Jason P Jue. All one needs to know about fog computing and related edge computing paradigms: A complete survey. *Journal of Systems Architecture*, 98:289–330, 2019.

[199] Charith Perera, Yongrui Qin, Julio C Estrella, Stephan Reiff-Marganiec, and Athanasios V Vasilakos. Fog computing for sustainable smart cities: A survey. *ACM Computing Surveys (CSUR)*, 50(3):1–43, 2017.

[200] Amir Vahid Dastjerdi, Harshit Gupta, Rodrigo N Calheiros, Soumya K Ghosh, and Rajkumar Buyya. Fog computing: Principles, architectures, and applications. In *Internet of things*, pages 61–75. Elsevier, 2016.

[201] Zeba Idrees, Zhuo Zou, and Lirong Zheng. Edge computing based iot architecture for low cost air pollution monitoring systems: a comprehensive system analysis, design considerations & development. *Sensors*, 18(9):3021, 2018.

[202] Rick Kazman, Mark Klein, and Paul Clements. *ATAM: Method for architecture evaluation*. Carnegie Mellon University, Software Engineering Institute Pittsburgh, PA, 2000.

[203] Diane Cook, Maureen Schmitter-Edgecombe, Aaron Crandall, Chad Sanders, and Brian Thomas. Collecting and disseminating smart home sensor data in the casas project. In *Proceedings of the CHI workshop on developing shared home behavior datasets to advance HCI and ubiquitous computing research*, pages 1–7. IEEE, 2009.

[204] Christian Montag and Harald Baumeister. Digital phenotyping and mobile sensing. *Digital Phenotyping and Mobile Sensing. 2nd ed. Springer International Publishing Cham, Switzerland*, 2023.

[205] Jimena Unzueta Saavedra, Emma A Deaso, Margot Austin, Laura Cadavid, Rachel Kraff, and Emma EM Knowles. Smartphone ecological momentary assessment and wearable activity tracking in pediatric depression: cohort study. *JMIR Formative Research*, 9:e66187, 2025.

[206] Cheryl Forchuk, Jonathan Serrato, Daniel Lizotte, Rupinder Mann, Gavin Taylor, and Sara Husni. Developing a smart home technology innovation for people with physical and mental health problems: considerations and recommendations. *JMIR mHealth and uHealth*, 10(4):e25116, 2022.

[207] Benjamin Agbo, Hussain Al-Aqrabi, Richard Hill, and Tariq Alsboui. Missing data imputation in the internet of things sensor networks. *Future Internet*, 14(5):143, 2022.

[208] World Health Organization. *Ethics and governance of artificial intelligence for health: large multi-modal models. WHO guidance*. World Health Organization, 2024.

[209] John Torous, Sandra Bucci, Imogen H Bell, Lars V Kessing, Maria Faurholt-Jepsen, Pauline Whelan, Andre F Carvalho, Matcheri Keshavan, Jake Linardon, and Joseph Firth. The growing field of digital psychiatry: current evidence and the future of apps, social media, chatbots, and virtual reality. *World psychiatry*, 20(3):318–335, 2021.

[210] Pasquale Bufano, Marco Laurino, Sara Said, Alessandro Tognetti, and Danilo Menicucci. Digital phenotyping for monitoring mental disorders: systematic review. *Journal of medical Internet research*, 25:e46778, 2023.

[211] Isaac Moshe, Yannik Terhorst, Kennedy Opoku Asare, Lasse Bosse Sander, Denzil Ferreira, Harald Baumeister, David C Mohr, and Laura Pulkki-Råback. Predicting symptoms of depression and anxiety using smartphone and wearable data. *Frontiers in psychiatry*, 12:625247, 2021.

[212] Thea Otte Andersen, Agnete Skovlund Dissing, Elin Rosenbek Severinsen, Andreas Kryger Jensen, Vi Thanh Pham, Tibor V Varga, and Naja Hulvej Rod. Predicting stress and depressive symptoms using high-resolution smartphone data and sleep behavior in danish adults. *Sleep*, 45(6):zsac067, 2022.

[213] Brunilda Balliu, Chris Douglas, Darsol Seok, Liat Shenhav, Yue Wu, Doxa Chatzopoulou, William Kaiser, Victor Chen, Jennifer Kim, Sandeep Deverasetty, et al. Personalized mood prediction from patterns of behavior collected with smartphones. *NPJ Digital Medicine*, 7(1):49, 2024.

[214] Amelia M Pellegrini, Emily J Huang, Patrick C Staples, Kamber L Hart, Jeanette M Lorme, Hannah E Brown, Roy H Perlis, and Jukka-Pekka J Onnela. Estimating longitudinal depressive symptoms from smartphone data in a transdiagnostic cohort. *Brain and Behavior*, 12(2):e02077, 2022.

[215] ML Tlachac, Michael V Heinz, Anastasia C Bryan, Arielle LaPreay, Geri Louise Dimas, Tingting Zhao, Nicholas C Jacobson, and Samuel S Ogden. Datasets of smartphone modalities for depression assessment: A scoping review. *IEEE transactions on affective computing*, 2025.

[216] ELISE L SHALOWITZ, PARDEEP S JHUND, MITCHELL A PSOTKA, Abhinav Sharma, MATTHEW G DIMOND, Trejeeve Martyn, Richard Nkulikiyinka, Mona Fiuzat, and DAVID P KAO. Where's the remote? failure to report clinical workflows in heart failure remote monitoring studies. *Journal of Cardiac Failure*, 2024.

[217] Marieke A Helmich. The duration-adjusted reliable change index: defining clinically relevant symptom changes of varying durations. *Assessment*, 31(7):1493–1507, 2024.

[218] Dimitra Georgiou, Stamatis Katsaounis, Panayiotis Tsanakas, Ilias Maglogiannis, and Parisis Gallos. Towards a secure cloud repository architecture for the continuous monitoring of patients with mental disorders. *Frontiers in Digital Health*, 7:1567702, 2025.

[219] Alexandru Rancea, Ionut Anghel, and Tudor Cioara. Edge computing in healthcare: Innovations, opportunities, and challenges. *Future internet*, 16(9):329, 2024.

[220] Health Level Seven International. Fhir overview. https://www.hl7.org/fhir/overview.html, 2024. Accessed: 2026-04-01.

[221] European Parliament and Council of the European Union. Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 (general data protection regulation). https://eur-lex.europa.eu/eli/reg/2016/679/oj, 2016. Accessed: 2026-04-01.

[222] U.S. Congress. Health insurance portability and accountability act of 1996 (hipaa), public law 104–191. https://www.govinfo.gov/content/pkg/PLAW-104publ191/pdf/PLAW-104publ191.pdf, 1996. Accessed: 2026-04-01.

[223] Banuchitra Suruliraj and Rita Orji. Federated learning framework for mobile sensing apps in mental health. In *2022 IEEE 10th International Conference on Serious Games and Applications for Health (SeGAH)*, pages 1–7. IEEE, 2022.

[224] Daphne James and Helen Warren-Forward. Research methods for formal consensus development. *Nurse researcher*, 22(3), 2015.

[225] Yuri Calleo, Simone Di Zio, and Francesco Pilla. Facilitating spatial consensus in complex future scenarios through real-time spatial delphi: A novel web-based open platform. *Futures & Foresight Science*, 5(3-4):e155, 2023.

[226] Ivan R Diamond, Robert C Grant, Brian M Feldman, Paul B Pencharz, Simon C Ling, Aideen M Moore, and Paul W Wales. Defining consensus: a systematic review recommends methodologic criteria for reporting of delphi studies. *Journal of clinical epidemiology*, 67(4):401–409, 2014.

[227] Prashant Nasa, Ravi Jain, and Deven Juneja. Delphi methodology in healthcare research: how to decide its appropriateness. *World journal of methodology*, 11(4):116, 2021.

[228] Caitlin Muhl, Kate Mulligan, Imaan Bayoumi, Rachelle Ashcroft, and Christina Godfrey. Establishing internationally accepted conceptual and operational definitions of social prescribing through expert consensus: a delphi study. *BMJ open*, 13(7):e070184, 2023.

[229] Zhida Shang. Use of delphi in health sciences research: a narrative review. *Medicine*, 102(7):e32829, 2023.

[230] William T Gattrell, Patricia Logullo, Esther J van Zuuren, Amy Price, Ellen L Hughes, Paul Blazey, Christopher C Winchester, David Tovey, Keith Goldman, Amrit Pali Hungin, et al. Accord (accurate consensus reporting document): A reporting guideline for consensus methods in biomedicine developed via a modified delphi. *PLoS Medicine*, 21(1):e1004326, 2024.

[231] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.

[232] Beenish M. Chaudhry. Daylio: Mood-quantification for a less stressful you. *mHealth*, 2(8), 2016.

[233] New Zealand Government. Privacy act 2020, 2020. Accessed: 2026-04-08.

[234] Eiko I Fried, Jessica K Flake, and Donald J Robinaugh. Revisiting the theoretical and methodological foundations of depression measurement. *Nature Reviews Psychology*, 1(6):358–368, 2022.

[235] Xuwei Shen. Data mining-based innovative model for mental health of college students using iot and big data analysis. *Soft Computing*, 27(19):14483–14495, 2023.

[236] Claes Wohlin, Per Runeson, Martin Höst, Magnus C Ohlsson, Björn Regnell, Anders Wesslén, et al. *Experimentation in software engineering*, volume 236. Springer, 2012.

Appendix A

Ethics Approval

This appendix contains the ethics approval documentation for this research.

Auckland University of Technology Ethics Committee (AUTEC)

9 December 2025

Minh Nguyen
Faculty of Design and Creative Technologies

Dear Minh

Re Ethics Application: **25/401 Expert Consensus on Safe and Personalised Systems for Depression Management**

Thank you for your responses to AUTEC's conditions.

Your ethics application has been approved for three years until 9 December 2028.

Standard Conditions of Approval

1. The research is to be undertaken in accordance with the Auckland University of Technology Code of Conduct for Research and as approved by AUTEC.
2. All public facing documents must have the AUTEC approval number and be of a high standard of spelling and grammar. Dates on the Information Sheet(s) and Consent Form(s) must be consistent.
3. Any amendments to the project must be approved by AUTEC prior to being implemented.
4. A progress report is due annually on the anniversary of the approval date.
5. A final report is due at the expiration of the approval period, or, upon completion of project.
6. Any serious or adverse events must be reported to AUTEC, this includes unforeseen issues that might affect continued ethical acceptability of the project.
7. AUTEC grants ethical approval only. You are responsible for obtaining management permission for access from any institution or organisation at which your research is being conducted and you need to meet all ethical, legal, public health, and locality obligations or requirements for the jurisdictions in which the research is being undertaken.

The application number and title need to be referenced on all correspondence related to this project.

All forms are available online http://www.aut.ac.nz/research/researchethics

For any enquiries, please contact the Secretariat at ethics@aut.ac.nz
(This is a computer-generated letter for which no signature is required)

The AUTEC Secretariat
Auckland University of Technology Ethics Committee

Cc: sanaz.zamani@autuni.ac.nz; sam.madanian@aut.ac.nz

Appendix B

Participant Information Sheet

This appendix contains the participant information sheet used for this research.

Participant Information Sheet – Experts

Date that data collection will start:

19 Sept 2025

Project Title

Expert Consensus on Safe and Personalised System for Depression Management

Kia Ora,

You are invited to participate in a research study on Internet of Things (IoT)-based safe and personalised depression management. This study is being conducted by Sanaz Zamani from Auckland University of Technology (AUT), under the supervision of Prof. Minh Nguyen and Dr. Samaneh Madanian (AUT), and Prof. Roopak Sinha (Deakin University). This study forms part of a PhD in Computer Science conducted by the main researcher.

What is the purpose of this research?

The primary aim of this study is to develop and validate a safe and personalised system for depression management using IoT technologies. This phase of the research uses the Delphi method, which is a structured process to gather expert opinions through multiple rounds of questionnaires to reach consensus on complex topics.

Your input will help identify essential system components, features, safety considerations, and technical or clinical requirements.

How was I identified and why am I being invited to participate in this research?

You are invited because you are a professional with expertise in one of the following areas:

- Clinicians: mental health professionals (e.g., psychiatrists, psychologists, general practitioners, or mental health nurses) with experience in depression care.
- Technical experts: professionals with expertise in IoT systems, cloud computing, system design, or AI/ML applications in healthcare.

Your expertise is critical for ensuring the system is safe, feasible, and effective.

How do I agree to participate in this research?

Participation is voluntary. Whether you choose to participate will neither advantage nor disadvantage you.

If you decide to take part, you will be asked to read this Information Sheet and sign the Consent Form. You may email the signed Consent Form back to the researcher.

Once your signed Consent Form is received, you will be provided with the link to the first questionnaire.

You may withdraw at any time prior to submitting your questionnaire responses by notifying the researcher, and your data will not be included in the study.

What will my participation involve?

- Providing your expert opinion on system design, safety, technical or clinical requirements, and practical considerations.
- Responding to questionnaires online over two rounds.
- Reviewing anonymised group responses in the second round and providing feedback or agreement.

A short scenario will be included in the questionnaire to illustrate how the system might work in practice.

What are the benefits?

To you as a participant:

- You may find the opportunity professionally rewarding by contributing your expert knowledge to the development of an emerging digital health system.
- You may choose to be acknowledged by name in resulting publications or presentations, recognising your contribution to the research.
- You will receive a summary of the study findings at completion, which may be of interest to your clinical or technical practice.

To the researcher:

- Your contribution will help create a validated, expert-informed framework for IoT-based depression management.

To the wider community:

- The findings may inform safer, personalised, and more effective digital mental health systems, ultimately benefiting patients and clinicians.

Acknowledgements

You may choose to be acknowledged by name in publications or presentations for your expert contribution. This is optional and will not affect your participation.

Will the results of the study be published?

The results of this research will be published in a Doctoral thesis. This thesis will be available to the general public through the AUT library. Results may be published in peer-reviewed, academic journals. Results will also be presented during conferences or seminars to wider professional and academic communities. You will not be identifiable in any publication.

Data Collection and Confidentiality

- Only your name, area of expertise, and email address will be collected to contact you for the second round.
- All survey responses are anonymous and analysed without identifying individuals.
- Data will be stored on AUT's secure OneDrive, encrypted and password-protected, accessible only to the supervisory team. Consent forms will be stored separately.
- No data will be stored on personal devices or non-approved cloud services. Both primary and backup storage are hosted in New Zealand.

What opportunity do I have to consider this invitation?

You will have up to two weeks to consider this invitation. You may ask questions or request further details before deciding.

Will I receive feedback on the results of this research?

A summary of the study findings will be provided to participants at the end of the study.

Who do I contact for further information about this research?

Please keep this Information Sheet and a copy of the Consent Form for your future reference. You are also able to contact the research team as follows:

Researcher Contact Details:

Sanaz Zamani (PhD Candidate at Auckland University of Technology)
Email Address: sanaz.zamani@autuni.ac.nz

Project Supervisor Contact Details:

Prof. Minh Nguyen (minh.nguyen@aut.ac.nz)
Dr. Samaneh Madanian (sam.madanian@aut.ac.nz)
Prof. Roopak Sinha (roopak.sinha@deakin.edu)

Approved by the Auckland University of Technology Ethics Committee on *09/12/2025*, AUTEC Reference number *25/401*.

Appendix C

Questionnaire

This appendix contains the questionnaire used in this research.

Questionnaires

Section 1: Questionnaire for Technical Experts

Scenario

Imagine you are tasked with designing and implementing an IoT-based intelligent system for personalised depression management. The system will operate in a smart-home environment for patients diagnosed with depression, aiming to continuously monitor risk factors and provide timely interventions while ensuring safety, personalisation, and reliability.

The proposed architecture integrates:

- Sensor layer: Non-invasive sensors (wearables, smart devices, and environmental sensors) to capture physiological (such as heart rate and sleep patterns), behavioural (such as activity levels and social interactions), and environmental data (such as light exposure and temperature).
- Edge computing: A gateway-based preprocessing for real-time analysis and reduced latency for critical alerts.
- Cloud infrastructure: Centralised processing for advanced analytics, machine learning model training, long-term data storage, and system health monitoring.
- User and clinician interface: Mobile app and dashboard for patients and healthcare professionals to receive alerts, summaries, and recommendations.
- Safety and personalisation features: Ensuring that every anomalies in patient behaviour will be detected, and personalised actions will be taken.

Your role is to evaluate the technical feasibility, reliability, safety and personalisation of this architecture, considering factors such as:

- Sensor integration and missing data handling.
- Edge vs. cloud trade-offs for latency and resource efficiency.
- AI-driven anomaly detection accuracy and personalisation strategies.
- Robustness of feedback loops and bi-directional communication.

The goal is to identify risks, bottlenecks, and optimisations that can make the system both technically sound and deployable in real-world settings.

A. Sensors and Data Collection

1. How appropriate do you consider the inclusion of three sensor categories (physiological, behavioural, and environmental) inside/outside smart homes for monitoring patients' behaviour?

a. Extremely Appropriate
b. Very Appropriate
c. Moderately Appropriate
d. Neutral
e. Somehow Appropriate
f. Slightly Appropriate
g. Not Appropriate at All

Please provide any additional feedback on your choice:

_____

B. System Initialisation

2. Rank replacement methods if real sensor data is unavailable

- Data imputation techniques
- Context-aware alternatives (daily questionnaires)
- Predictive modelling
- Rule-based input (clinician-provided)

3. How effective do you think replacement methods (imputation, questionnaires, predictive models) are for filling missing data gaps?

a. Extremely Effective
b. Very Effective
c. Moderately Effective
d. Neutral
e. Somehow Effective
f. Slightly Effective
g. Not Effective at All

Please provide any additional feedback on your choice:

_____

4. List key parameters to calibrate during system setup to reduce false positives.

_____

5. Which approach is more appropriate for model initialisation?

a. user-specific model initialisation (personalised models)
b. generic pre-trained models

Please provide any additional feedback on your choice:

_____

C. Edge-Cloud Data Processing

6. Edge-based analysis reduces network load and accelerates emergency alerts.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

7. Periodic cloud retraining adapts to individual behaviour and reduces false alarms.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

8. Edge processing (feature extraction/classification) is feasible on current IoT gateways.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

9. Discuss trade-offs between edge real-time processing and cloud post-processing.

_____

10. Are you aware of specific cloud services or platforms (such as Azure, AWS, Google Cloud) that could support the design, deployment, or operation of the proposed IoT-based depression management system?
a. Yes – Please specify: _____
b. No

11. Please indicate which cloud services or features you consider most relevant or beneficial for this system (select all that apply):
   a. IoT device management (such as Azure IoT Hub, AWS IoT Core)
   b. Real-time data processing and analytics (such as AWS Lambda, Azure Stream Analytics)
   c. Machine learning model training and deployment (such as Azure ML, Amazon SageMaker)
   d. Data storage and database services (such as AWS S3, Azure Cosmos DB)
   e. Security and identity management (such as AWS IAM, Azure Active Directory)
   f. Cloud edge computing integration (such as AWS Greengrass, Azure IoT Edge)
   g. Other (please specify): _____

D. Monitoring Mechanism

12. What is the optimal frequency for system health checks (battery/connectivity)?
   a. Daily
   b. Every Other Day
   c. Twice a week
   d. Weekly
   e. Fortnightly
   f. Monthly
   g. Never

E. Anomaly Detection

13. Which machine learning techniques do you consider best for detecting behavioural anomalies?

_____

14. Edge-based anomaly detection reduces latency.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

   _____

15. Post-processing (e.g., filtering) minimises false alarms.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

   _____

F. Feedback Process

16. User feedback improves long-term detection accuracy.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

17. Automatic logging/cloud syncing of emergencies is crucial for audit and retraining.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

18. Bidirectional communication (clinician updates to devices) is necessary.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

G. Technical Practice Context and Perceptions

19. My current work in IoT/smart-home or health informatics involves significant technical challenges in designing or implementing reliable healthcare monitoring systems.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

20. I actively track specific performance metrics or "red flags" (such as latency spikes, packet loss, battery depletion) that can impact the reliability of IoT healthcare systems.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

21. I am highly confident in using and integrating advanced IoT, AI, and cloud-edge technologies in healthcare system development.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

22. IoT-based smart-home systems have strong potential to enhance depression monitoring and intervention compared to conventional healthcare solutions.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

23. Lack of adequate data quality, interoperability standards, or security measures is a major challenge in IoT-based healthcare system development.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

24. There are unmet technical requirements in current healthcare systems that could be addressed by more advanced IoT-enabled monitoring and intervention platforms.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

25. The proposed IoT-based depression management system is technically feasible and could operate effectively in real-world settings.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

26. Are you aware of any existing IoT-based or digital monitoring systems designed for depression or mental health management?
a. Yes – Please specify: _____
b. No

H. Legal and Regulatory Considerations

27. Are there particular regulations or standards that must be followed when designing IoT-based healthcare monitoring systems?
(Please consider areas such as data privacy, interoperability standards, and medical certifications)

_____

28. What legal or compliance challenges do you foresee in implementing such systems? (such as data security, device certification, interoperability, patient consent and data ownership)?

_____

29. How should regulatory frameworks adapt depending on whether the system is used primarily as:
a. A patient self-management tool?
b. A clinician-facing decision support system?

(Please consider differences in requirements for approval, data governance, liability, and clinical oversight.)

_____

30. From your technical expertise, please suggest any additional questions or areas of inquiry that you believe are important to include in this questionnaire to better evaluate or fine-tune the proposed IoT-based depression management system architecture.

_____

31. Please share any additional comments, concerns, or ideas related to the above questions or the proposed IoT-based depression management system.
(You may address technical feasibility, ethical considerations, security risks, or potential improvements)

_____

32. How important is compliance with interoperability standards (such as HL7 and FHIR) for ensuring the system integrates effectively with healthcare platforms (such as EHRs)?

a. Extremely Important
b. Very Important
c. Moderately Important
d. Neutral
e. Somehow Important
f. Slightly Important
g. Not Important at All

Please provide any additional feedback on your choice:

_____

33. How important is adherence to recognised cybersecurity frameworks (such as NIST, ISO 27001) for ensuring patient data protection and system reliability?

a. Extremely Important
b. Very Important
c. Moderately Important
d. Neutral
e. Slightly Important
f. Not Important at All

Please provide any additional feedback on your choice:

_____

34. What additional security or compliance measures do you recommend for an IoT-based healthcare system?

_____

Section 1: Questionnaire for Healthcare Professionals

Scenario:

Imagine you are caring for an adult patient diagnosed with moderate depression, who is living independently at home. The patient has been prescribed medication and receives periodic therapy sessions but often struggles with daily routine and mood fluctuations.

To support this patient, an Internet of Things (IoT)-based depression management system is introduced. The system uses non-invasive sensors (such as smart devices, wearables, and environmental sensors) and Artificial Intelligence (AI) algorithms to monitor physiological (such as heart rate and sleep measurements), behavioural (such as social interactions and activity levels), and environmental factors (such as light exposure and temperature).

When the system detects anomalies suggesting a potential worsening condition, it can:

- Send alerts (to the patient, GP, or specialist, depending on severity).
- Provide personalised recommendations (such as reminders and coping strategies).
- Enable real-time clinician access to summaries and alerts through an application or dashboard.

Your role in this questionnaire is to evaluate how appropriate and useful these features are in clinical practice, identify potential risks, and suggest improvements to make the system clinically relevant and safe.

A. Data Collection

1. How important do you think physiological, behavioural, and environmental data are for detecting early signs of anomalies in users?

a. Extremely Important
b. Very Important
c. Moderately Important
d. Neutral
e. Somehow Important
f. Slightly Important
g. Not Important at All

Please provide any additional feedback on your choice:

_____

2. Rate the potential impact of each data type on emotional well-being and depression management. The impact refers to how strongly each data type can influence the ability to detect mood changes or depression-related patterns and support timely intervention.

Data	Very Low Impact	Low Impact	Slightly Low Impact	No impact	Slightly High Impact	High Impact	Very High Impact
Physiological (e.g., heart rate)							
Behavioural (e.g., social connections)							
Environmental (e.g., weather conditions)							

Please provide any additional feedback on your choice:

_____

3. Personalised recommendations and actions generated by the system (tailored to individual characteristics and sensitivities) improve the quality of depression management.

a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

4. How useful do you think physiological, behavioural, and environmental data are for making personalised decisions for users?
   a. Extremely Useful
   b. Very Useful
   c. Moderately Useful
   d. Neutral
   e. Somehow Useful
   f. Slightly Useful
   g. Not Useful at All

   Please provide any additional feedback on your choice:

   _____

B. System Initialisation

5. Assign the most suitable healthcare professional for each service:

Service	General Practitioner	Mental Health Practitioner
Determining patient sensitivity		
Daily check-ins		
Receiving urgent alerts		
Providing immediate care		
Providing advanced care		

Please provide any additional feedback on your choice:

_____

6. Who should be responsible for initial system setup?
   a. General Practitioner
   b. Mental Health Practitioner
   c. Patients
   d. System's Technical Support
   e. Others

   Please provide any additional feedback on your choice:

   _____

7. Involving patients in system setup helps personalise the environment rather than causing inconvenience.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

   _____

C. Data Processing

8. Delayed alerts (due to technological issues) reduce confidence in the system.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

D. Monitoring Mechanism

9. Map useful alert levels to the symptom severity:

Severity	Local Audio Alerts	Alert to GP	Alert to Specialist
Severe			
Moderate			
Mild			

Please provide any additional feedback on your choice:

_____

10. Local alert mechanisms are useful for moderate and mild conditions.
a. Strongly Agree
b. Agree
c. Somehow Agree
d. Neutral
e. Somehow Disagree
f. Disagree
g. Strongly Disagree

Please provide any additional feedback on your choice:

_____

11. What is your preferred method of receiving alerts?
a. App notifications
b. SMS
c. Integrated EMR alerts

Please provide any additional feedback on your choice:

_____

E. Anomaly Detection

12. How confident would you be in trusting an emergency alert if the system achieves >= 90% accuracy in real-world conditions?
   a. Extremely Confident
   b. Very Confident
   c. Moderately Confident
   d. Neutral
   e. Somehow Confident
   f. Slightly Confident
   g. Not Confident at All

   Please provide any additional feedback on your choice:

   _____

13. Frequent false positives reduce clinical trust.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

   _____

14. In some cases, missing an emergency alert (false negative) may be acceptable if it reduces unnecessary alerts and interventions.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

   _____

F. Feedback Process

15. Users and caregivers should have access to event summaries through an app or portal.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

---

16. Allowing patients to confirm or cancel false alarms reduces unnecessary interventions.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

G. Clinical Practice Context and Perceptions

17. My current clinical work involves significant challenges in identifying and addressing patients' needs related to depression.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

---

18. I actively monitor for specific "red flags" or early warning signs that indicate worsening depression or increased risk in my patients.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

---

19. How would you describe your own level of digital literacy?
   a. Very Low – I am not confident and require significant assistance with digital tools.
   b. Low – I require regular assistance with most digital tools.
   c. Slightly Low – I can use only basic digital tools and often need help with intermediate features.
   d. Moderate – I can use basic digital tools but may require support for advanced features.
   e. Slightly High – I can use most digital tools but may need occasional help with advanced technologies.
   f. High – I can use most technologies without difficulty.
   g. Very High – I am highly confident and can independently use and learn new technologies.

   Please provide any additional feedback on your choice:

   _____

20. Digital health technologies (such as IoT devices, smart-home monitoring) can significantly improve the quality and timeliness of mental health care.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

   _____

21. Lack of appropriate tools and systems is a major challenge in delivering effective depression care.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

   _____

22. I believe there are unmet needs in my clinical practice that could be addressed by technology-assisted monitoring and intervention systems.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

_____

23. The proposed IoT-based depression management system could improve depression care in my clinical setting.
   a. Strongly Agree
   b. Agree
   c. Somehow Agree
   d. Neutral
   e. Somehow Disagree
   f. Disagree
   g. Strongly Disagree

   Please provide any additional feedback on your choice:

   _____

24. Are you aware of any existing digital apps or monitoring systems for depression management (such as IoT-based systems, mobile applications, or clinical platforms)?
   a. Yes – Please specify: _____
   b. No

H. Legal and Regulatory Considerations

25. Please indicate any specific rules, regulations, or legislation that should be considered when designing the depression management system (such as data privacy, medical device regulations, patient consent laws).

   _____

26. What legal or ethical considerations must be addressed when implementing such a system in clinical practice?

   _____

27. From your perspective, how do regulatory requirements differ when the system is used primarily as:
   c. A support tool for patients themselves?
   d. A clinical decision support tool for healthcare professionals?

   Please provide any additional feedback on your choice:

   _____

28. Please share any additional comments or suggestions related to the above questions or the overall topic of depression management and monitoring systems.

   _____

Appendix D

Participant Feedback Reports

This appendix contains the feedback reports provided to participants following the questionnaire.

1. Expert Questionnaire Results – Technical Experts

A structured questionnaire was administered to domain experts (14 experts) to evaluate the proposed IoT-based depression management system. A consensus-based analysis approach was used, with a 70% agreement threshold considered sufficient to indicate expert support. Agreement levels above 85% were interpreted as strong consensus, while 100% indicated full unanimity.

Overall, the results demonstrate substantial expert endorsement of the proposed architecture, particularly in relation to multimodal sensing, hybrid edge-cloud processing, and personalised modelling. These findings suggest that the system is technically feasible according to expert opinion and aligns with current expectations for healthcare IoT solutions. However, some lower-consensus items highlight important design sensitivities that require careful consideration.

1.1. Consensus summary of closed-ended questions

Table 1 shows the consensus percentage on different components of the architecture, collected by closed-ended questions.

Table 1: Consensus summary – technical experts

Component	Consensus (%)
Multimodal sensing	92.9
Missing data handling	78.6
Edge-based analysis	100
Cloud retraining	100
Edge feasibility	85.7
Latency reduction (edge)	85.7
Post-processing	92.9
User feedback	100
Cloud logging	100
Bidirectional communication	92.9
Technical feasibility	85.7

The following subsections provide details on the closed-ended questions.

1.1.1. Multimodal sensing and data handling

Experts strongly supported the inclusion of physiological, behavioural, and environmental data (92.9%). The methods for handling missing data achieved 78.6% consensus, indicating acceptance but highlighting the need for careful implementation. Ranking results showed a preference for context-aware approaches (e.g., questionnaires) followed by predictive modelling, data imputation techniques, and rule-based input.

These findings suggest that system design should prioritise adaptive, context-sensitive data handling mechanisms, particularly in dynamic real-world environments.

1.1.2. Edge-cloud architecture and real-time monitoring

There was unanimous agreement (100%) that edge-based processing reduces network load and accelerates emergency alerts, while 85.7% agreed that it reduces latency. Additional components, including edge feasibility (85.7%) and post-processing (92.9%), were also strongly supported. This validates a hybrid architecture in which the edge handles real-time detection and the cloud supports long-term learning and optimisation.

1.2. Thematic summary of open-ended responses

Table 2 shows the key insights extracted from open-ended responses for each theme related to the proposed architecture.

Table 2: Key themes from qualitative analysis – technical experts

Theme	Key insight
Personalised baseline	Individual patterns must guide detection
Sensor calibration	Accuracy, noise, and placement are critical
Context-aware rules	Behaviour must be interpreted within context
Multi-sensor fusion	Combining modalities improves robustness
Thresholds & smoothing	Alerts should depend on sustained changes
Missing data tolerance	The system must handle incomplete data robustly
User burden control	Over-prompting can affect behaviour

These themes reinforce that personalisation, calibration, and context-awareness are central to reducing false positives and improving reliability. The following subsections explain the details of the insights based on open-ended questions.

1.2.1. Model initialisation and anomaly detection

Experts favoured user-specific models, with some support for hybrid approaches starting from pre-trained models. Recommended techniques included unsupervised and semi-supervised methods (e.g., autoencoders, Isolation Forest, LSTM), reflecting the limited availability of labelled data.

1.2.2. Cloud services, operations, and governance

IoT device management was the most frequently selected cloud capability. Daily system health checks were preferred, highlighting the need for continuous monitoring. Experts also emphasised compliance with data privacy, interoperability, and medical standards (e.g., GDPR, HIPAA, HL7/FHIR, ISO).

1.3. Overall interpretation

The results demonstrate broad expert support for the proposed system, particularly its emphasis on personalisation, hybrid edge-cloud processing, and multimodal sensing. However, lower agreement in some areas indicates the need for careful handling of time-sensitive and safety-critical alerts and system assumptions.

These findings directly inform the design of the proposed architecture by validating key decisions related to adaptive modelling, real-time processing, and system reliability.

2. Expert Questionnaire Results – Clinical Experts

A questionnaire was administered to eight clinical experts to evaluate the proposed IoT-based depression management system from the perspectives of clinical usefulness, alert handling, workflow integration, and regulatory considerations. A 70% agreement threshold was used to indicate consensus, with ≥85% representing strong agreement and 100% full unanimity.

Overall, the results demonstrate strong support for multimodal data integration, personalised decision-making, and reliable alerting. At the same time, experts highlighted critical concerns regarding alert reliability, missed critical events, patient engagement, and regulatory compliance.

2.1. Consensus summary of closed-ended questions

Table 3 summarises the level of agreement across key components. All items achieved consensus except for the acceptance of missed emergency alerts (25%), indicating strong expert rejection of false negatives in safety-critical contexts.

Table 3: Consensus summary - clinical experts

Components	Consensus (%)
Multimodal data for anomaly detection	100
Personalised recommendations improve care	100
Multimodal data supports personalised decisions	100
Patient involvement in setup	100
Delayed alerts reduce trust	100
False positives reduce clinical trust	100
Local alerts for mild/moderate cases	75
Trust in alerts at ≥90% accuracy	75
Event summaries via app/portal	75
Patients confirm/cancel false alarms	87.5
Missing an alert may be acceptable	25
Digital literacy	100
Digital health can improve care	100
Lack of tools is a major challenge	100
Unmet needs could be addressed by technology	75
The proposed system could improve care	75

The following subsections provide more details on responses from closed-ended questions in the questionnaire for clinical experts.

2.1.1. Multimodal data and personalisation

Experts unanimously supported the use of physiological, behavioural, and environmental data for early anomaly detection and personalised decision-making. However, ranking results indicate that physiological and environmental data were perceived as most influential, while behavioural data received comparatively lower influence scores. This suggests that although multimodal sensing is essential, not all data types contribute equally to clinical insight.

2.1.2. Alerting, trust, and safety

There was broad agreement that delayed alerts and frequent false positives erode trust, underscoring the importance of reliable, timely system responses. Moderate consensus (75-87.5%) supported local alerts, high-confidence alerting, event summaries, and patient confirmation of false alarms, indicating preference for user-in-the-loop mechanisms.

A key finding is the strong reluctance to tolerate missed emergency alerts (25%), demonstrating that experts prioritise sensitivity and safety over reducing alert frequency in critical scenarios.

Alert delivery preferences were distributed across app notifications, EMR integration, and SMS, suggesting the need for multi-channel alert dissemination. In addition, experts showed consistent mapping between alert levels and symptom severity (severe, moderate, mild), supporting structured alert hierarchies within the system design.

2.1.3. Clinical workflow and governance

Role allocation revealed a clear division of responsibilities: mental health practitioners were preferred for sensitivity assessment and advanced care, while general practitioners were more frequently associated with urgent alerts and immediate intervention. Initial system setup was also most commonly assigned to mental health practitioners, supporting a layered, role-specific care model.

Despite unanimous agreement that lack of tools is a major challenge (100%), fewer experts reported active monitoring practices (50%) or significant clinical challenges (37.5%), suggesting variability in current workflows and potential gaps between perceived needs and practice.

Awareness of existing digital depression tools was limited, with most respondents who answered this item being unfamiliar with such systems, reinforcing the novelty and potential value of the proposed approach.

Experts consistently emphasised regulatory requirements, including data privacy (e.g., Privacy Act 2020), informed consent, medical device compliance, clinical safety, cybersecurity, and accessibility. Requirements varied by context, with patient-facing tools focusing on safe self-management and clinical systems requiring stricter validation, accountability, and auditability.

2.2. Thematic summary of open-ended responses

Table 4 summarises the key themes identified from open-ended responses.

Table 4: Main themes from open-ended responses – clinical experts

Theme	Key insight
Multimodal sensing	Use combined physiological, behavioural, and environmental data
Personalisation	Tailor recommendations and thresholds to the individual
Trustworthy alerts	Reduce delay and false positives; keep alerts clinically reliable
User-in-the-loop design	Allow patients to confirm or cancel false alarms
Role-specific workflows	Distribute responsibilities across clinical roles
Regulatory compliance	Address privacy, consent, safety, and medical device rules
Engagement and adherence	Minimise burden to support long-term use

Experts consistently emphasised the importance of privacy, informed consent, medical device regulation, clinical safety, cybersecurity, and accessibility. Regulatory expectations varied depending on system use: patient-facing tools were associated with privacy and safe self-management, while clinical deployment required stricter validation, accountability, and compliance.

A notable insight relates to patient engagement. Several experts highlighted that individuals with severe depression may struggle with adherence and active participation, indicating that system design must minimise user burden and support passive monitoring where possible.

Furthermore, limited awareness of existing digital depression tools among participants suggests that such systems are not yet widely integrated into routine clinical practice, reinforcing the potential novelty and relevance of the proposed solution.

2.3. Overall interpretation

Using the 70% consensus threshold, most components of the proposed system were supported, with several achieving full unanimity. The findings validate the importance of multimodal data integration, personalised interventions, and reliable alerting mechanisms.

Crucially, the results demonstrate that missed critical alerts are largely unacceptable, reinforcing the need for high-sensitivity detection. At the same time, variability in current clinical practices, limited awareness of existing tools, and concerns around patient engagement highlight important real-world implementation challenges.

Overall, the proposed system is supported as a clinically relevant and promising approach, provided that reliability, regulatory compliance, and user engagement are carefully addressed.