RASF-LCS: Ranked Attribute Selection and Distance-Based Rule Filtering for Interpretable Credit Scoring

Mohamed Azhar Ahamed Zahvie
Wellington Institute of Technology
Wellington, New Zealand
zahvie@gmail.com

Abubakar Siddique
Auckland University of Technology
Auckland, New Zealand
abubakar.siddique@aut.ac.nz

Trung Nguyen
Auckland University of Technology
Auckland, New Zealand
trung.nguyen@aut.ac.nz

Muhammad Iqbal
Higher Colleges of Technology
Fujairah, United Arab Emirates
miqbal1@hct.ac.ae

Will N. Browne
Queensland University of Technology
Brisbane, Australia
will.browne@qut.edu.au

Abstract

Machine learning is widely used in credit scoring, but many high-performing models lack interpretability, limiting their use in regulated domains. Rule-based approaches such as Learning Classifier Systems (LCS) offer a balance between accuracy and explainability. This paper introduces Ranked Attribute Selection with Midpoint Filtering (RASF), an extension to LCS that enhances feature selection and rule validation. RASF combines mutual information-based feature ranking, guided attribute selection, and distance-based rule filtering. Experiments on loan approval datasets show that RASF improves accuracy by 3–5% over standard LCS while maintaining interpretable, rule-based outputs. These results highlight the potential of RASF-LCS for explainable credit decision-making.

CCS Concepts

• **Computing methodologies → Rule learning**; *Supervised learning by classification.*

Keywords

Credit approval, learning classifier systems, interpretability, feature selection, mutual information, explainable AI

ACM Reference Format:
Mohamed Azhar Ahamed Zahvie, Abubakar Siddique, Trung Nguyen, Muhammad Iqbal, and Will N. Browne. 2026. RASF-LCS: Ranked Attribute Selection and Distance-Based Rule Filtering for Interpretable Credit Scoring. In *Genetic and Evolutionary Computation Conference (GECCO Companion '26), July 13–17, 2026, San Jose, Costa Rica.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3795101.3805344

1 Introduction

Credit scoring in banking aims to predict borrower risk and typically relies on statistical models such as logistic regression [1]. In recent years, complex machine learning (ML) models (e.g., neural networks, ensembles) have been applied to credit data, yielding higher accuracy [3, 9]. However, these models often operate as "black boxes" [13], creating challenges for transparency and fairness [5]. Regulatory frameworks (e.g., Equal Credit Opportunity Act, GDPR) require that automated lending decisions be explainable [5, 6]. Thus, there is a strong demand for models that balance accuracy with explainability.

Rule-based ML methods provide one explainable AI solution for credit scoring: they represent knowledge as human-readable IF–THEN rules. A prominent example is the Learning Classifier System (LCS) [16, 17, 20], which evolves a population of IF–THEN rules using genetic operators. LCS models naturally yield interpretable rule lists and have been applied in classification and reinforcement learning contexts [4, 14, 18, 21]. However, standard LCS implementations treat all features equally and may generate overly general or irrelevant rules [21]. In credit risk, where some features (e.g., income, credit history) are known to be crucial, guiding the model with feature importance can improve rule relevance and decision accuracy.

In this work, we address these limitations by proposing a Ranked Attribute Selection with Midpoint Filtering (RASF) framework that enhances the Learning Classifier System (LCS) covering mechanism for credit approval tasks. The proposed RASF-LCS introduces three main contributions. First, mutual information (MI) feature ranking is used to identify informative attributes. Second, a rank-guided randomised attribute selection strategy biases rule construction toward high-MI features while preserving exploration of lower-ranked ones. Third, a midpoint-based distance filtering mechanism evaluates newly generated rules using Euclidean distance to ensure contextual relevance before acceptance.

These components are integrated into the LCS covering phase, enabling rules to be both informative and instance-relevant. We evaluate RASF-LCS on two public credit datasets, comparing against a baseline LCS and logistic regression. Results show that RASF improves predictive performance while maintaining interpretable, rule-based outputs, supporting a balance between accuracy and transparency in credit decision-making.

2 Literature Review

Conventional credit scoring approaches rely on statistical techniques such as logistic regression and discriminant analysis [1].

GECCO Companion '26, San Jose, Costa Rica

ACM ISBN 979-8-4007-2488-6/2026/07
https://doi.org/10.1145/3795101.3805344

These models are valued for their transparency but are often outperformed by modern machine learning (ML) methods. Neural networks and related approaches have demonstrated improved accuracy [3, 8, 9], but at the cost of reduced interpretability. In regulated financial domains, where decisions must be justified, this trade-off between predictive performance and interpretability remains a central challenge.

Over the past two decades, ML techniques have been widely adopted for credit risk prediction, including deep learning and ensemble methods [8]. Ensemble approaches such as random forests and gradient boosting improve generalisation and can yield accuracy gains of 2% to 15%. However, these gains come with increased complexity, as models with many parameters are difficult to interpret, limiting their use in high-stakes applications.

Explainable AI (XAI) seeks to address this limitation through post-hoc explanations and inherently interpretable models [13]. Post-hoc methods (e.g., LIME, SHAP) approximate black-box behaviour but may not faithfully capture decision logic in complex settings. In contrast, interpretable models such as decision trees, rule lists, and generalised additive models provide transparency by design, though they may sacrifice performance or become difficult to interpret when overly complex.

Learning Classifier Systems (LCS) are evolutionary, rule-based machine learning algorithms that construct populations of IF–THEN rules through reinforcement learning and genetic algorithms [20]. They provide a balance between interpretability and adaptability by decomposing complex decision spaces into human-readable rules. Recent work has extended LCS with richer representations and improved scalability. In particular, code-fragment-based LCS approaches introduce reusable building blocks, improving generalisation and learning efficiency [2, 11, 12]. In parallel, LCS has been explored as an interpretable AI framework for decision-making in complex domains [15].

Despite these advantages, its application to credit scoring remains limited [10]. A study on XCSF introduces a guided mutation operator using accuracy-weighted covariance of stored samples, to improve learning efficiency and scalability in high-dimensional problems [19]. However, while guided XCSF focuses on optimising classifier shape for regression using mutation only, RASF-LCS instead emphasises feature ranking and rule filtering to enhance interpretability and performance in the covering process.

Feature selection addresses this issue by identifying informative variables. Mutual information (MI) is widely used as a filter-based criterion due to its ability to capture both linear and non-linear relationships [7]. In credit scoring, MI has been applied to identify key predictors [8], yet it has not been systematically integrated into LCS, where feature selection is often implicit or manually specified.

3 Methods

We enhance the *ExSTraCS* framework [20] by integrating a Ranked Attribute Selection and Midpoint Filtering (RASF) process into the rule creation phase (Figure 1). The improvement focuses on the covering process, where new rules are generated. In RASF-LCS, covering involves: (1) computing feature importance and rankings; (2) selecting attributes via Algorithm 1; (3) generating rule conditions from instance values; and (4) applying Algorithm 2 to retain or discard rules. Accepted rules are added to the population and match set, while subsequent LCS processes (parameter updates, subsumption, and GA) remain unchanged.

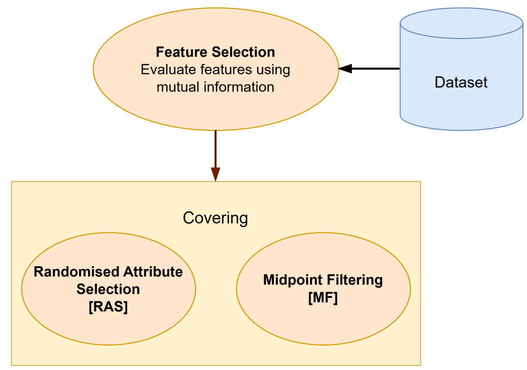

Figure 1: Overview of RASF-LCS.

3.1 Feature Importance via Mutual Information

We rank input attributes based on their mutual information (MI) with the output attribute. Given the training data (X, y), where each instance has attributes X_j and class label y, the mutual information $I(X_j; y)$ quantifies the reduction in uncertainty about y provided by X_j. It is defined as $I(X_j; Y) = \sum_{x,y} p(x, y) \log \frac{p(x,y)}{p(x)\,p(y)}$, where $p(x, y)$ represents the joint probability of the feature value x and class y, while $p(x)$ and $p(y)$ are the respective marginal probabilities. Features that exhibit higher mutual information (MI) scores are more predictive of the class. We calculate $I(X_j; Y)$ for each feature using scikit-learn and rank the features in descending order of I. A certain percentage (e.g., $p\%$) of the top features is then selected for rule generation. This ranked list helps the LCS prioritise the most pertinent attributes initially.

3.2 Rank-Guided Randomised Attribute Selection

For each new instance, the LCS selects up to R attributes to form a rule condition using a rank-guided randomised selection strategy (Algorithm 1). Top-ranked features (based on importance) are first shuffled, and an initial candidate set of size R is formed. From this set, $m = \lceil pR \rceil$ features are randomly selected, while the remaining $R - m$ features are drawn from lower-ranked attributes. The combined set of R features defines the rule condition. This approach balances exploitation of high-importance features with exploration of lower-ranked ones, improving diversity while preserving relevance.

3.3 Midpoint-Based Rule Filtering

Once a candidate rule is formed, a contextual relevance check is applied before adding it to the LCS population (Algorithm 2). The process is as follows:

- **Midpoint calculation:** For each selected attribute i, compute the midpoint m_i. For continuous attributes with bounds $[l_i, u_i]$, $m_i = (l_i + u_i)/2$. For categorical attributes, $m_i = x_i/2$.

Algorithm 1 Rank-guided randomised attribute (feature) selection.

```
Input: Ranked feature list F, rule size R, fraction p
Output: Selected feature subset S for rule generation
Copy and shuffle top of F into list F′
Take first R features from F′ as candidate set C
Compute m = ⌈p R⌉, n = R − m
Randomly select m features from C as S_1
Randomly select n features from remaining lower-ranked features as S_2
S ← S_1 ∪ S_2
return S
```

- **Distance computation:** Calculate the Euclidean distance between the midpoint vector $\mathbf{m}$ and the instance $\mathbf{x}$:

$$d = \sqrt{\sum_{i \in S} (m_i - x_i)^2}.$$

- **Filtering:** If $d \leq \theta$, accept the rule; otherwise, reject it and update the threshold $\theta \leftarrow d + 0.1d$.

This adaptive filtering ensures that only rules sufficiently close to the instance are retained, improving contextual relevance. The covering process repeats until enough rules are generated under the updated threshold.

Algorithm 2 Midpoint-based distance filtering

```
Input: Current instance x, new rule with condition bounds, initial threshold θ
Output: Accept or reject the rule
for each attribute i in the rule do
    if attribute i is continuous then
        m_i ← (l_i + u_i)/2 {midpoint of bounds}
    else
        m_i ← x_i/2 {for discrete attribute}
    end if
end for
Compute d = sqrt(Σ_i (m_i − x_i)^2)
if d < θ then
    Accept rule (add to population and match set)
else
    θ ← d + 0.1 d {raise threshold}
    Reject rule
end if
```

4 Experimental Setup and Data

4.1 Data

This study uses two credit datasets: a *Loan Approval Classification* dataset (45,000 records, 14 attributes) and a *Loan Approval Prediction* dataset (58,645 records). Both include numerical and categorical features such as age, income, employment, loan characteristics, and credit history, with a binary target indicating approval status.

The primary dataset is imbalanced (10,000 approvals vs. 35,000 rejections), reflecting realistic financial decision scenarios. Features were standardised using attribute IDs (A01–A14), with details summarised in Table 1.

Table 1: Loan Approval Classification Dataset.

ID	Feature	Type	Description
A01	Age	Continuous	Applicant's age
A02	Income	Continuous	Applicant's monthly or annual income
A03	Work Experience	Continuous	Years of employment
A04	Loan Amount	Continuous	Requested loan value
A05	Loan Purpose	Categorical	Type of loan (e.g., personal, education)
A06	Home Ownership	Binary	1 = Own home, 0 = Rent/Other
A07	Loan-to-Income Ratio	Continuous	Ratio of loan amount to income
A08	Interest Rate	Continuous	Applied loan interest rate
A09	Credit Score	Continuous	Normalised credit-worthiness indicator
A10	Loan Grade	Categorical	Assigned grade from credit bureau
A11	Education	Categorical	Highest qualification level
A12	Employment Type	Categorical	Full-time, self-employed, etc.
A13	Previous Default	Binary	1 = Default history, 0 = No default
A14	Approval Status	Binary	Target class: 1 = Approved, 0 = Rejected

Table 2: Test Accuracy.

Dataset	LCS	RASF-LCS	Logistic Regression
Loan Approval	81.87%	86.51%	88.94%
Loan Approval Prediction	79.80%	83.40%	90.02%

A consistent preprocessing pipeline was applied:

- **Missing values:** Mean (numerical) and mode (categorical) imputation.
- **Categorical encoding:** One-hot and label encoding to avoid ordinal bias.
- **Feature scaling:** Normalisation to [0,1] for numerical attributes.
- **Class distribution:** No resampling applied to preserve real-world imbalance.

4.2 RASF-enhanced ExSTraCS

The proposed RASF mechanism introduces adaptive midpoint-based distance filtering, where the threshold increases incrementally (10%) until a suitable rule is generated. Attribute selection follows a hybrid strategy: half of the features are chosen from top-ranked (mutual information) attributes, while the remainder are sampled from lower-ranked features to balance relevance and diversity. The RASF-ExSTraCS used the parameter values that have been commonly used in the literature [20**?**] and the developed model was trained using 100,000 iterations.

Model performance was evaluated across predictive accuracy, efficiency, and interpretability. Interpretability was assessed using:

- **Rule compactness:** Fewer and shorter rules indicate better interpretability.
- **Feature usage frequency:** Identifies key attributes influencing decisions.
- **Rule transparency:** Human-readable IF–THEN rules support traceability, critical for credit decision contexts.

5 Experimental Results

We evaluated the RASF-LCS on two publicly available credit datasets ("Loan Approval" and "Loan Approval Prediction"). Both contain applicant financial attributes and a binary outcome (approve/deny).

Each LCS model (baseline and RASF) was run for 100,000 iterations over 10 random train/test splits. We compare test accuracies averaged over runs. A standard logistic regression (LR) model is also used as a benchmark.

Table 2 summarises the accuracy of the baseline LCS, logistic regression, and RASF-LCS. In both datasets, RASF-LCS achieves higher accuracy. For example, on the Loan Approval data, test accuracy improved from 81.9% (base LCS) to 86.5% (RASF-LCS). The LR model obtains slightly higher accuracy (around 89%), but without producing explicit rules.

In addition to accuracy, we evaluated interpretability through the learned rule population. The RASF-LCS produced approximately 6.5k IF–THEN rules, with key features such as loan-to-income ratio, default history, interest rate, and home ownership appearing most frequently—consistent with known credit risk factors. Less relevant attributes (e.g., education, gender) were rarely used.

The rules are compact (typically 1–3 conditions) and easy to interpret. Examples include:

- *Approve:* If loan-to-income < 0.35 and no prior default.
- *Reject:* If previous default = Yes and loan-to-income ≥ 0.5.
- *Approve:* If home ownership $\in$ {OWN, MORTGAGE}, loan-to-income < 0.4, and interest rate < 0.12.

These rules reflect intuitive credit logic, where high debt burden or default history increases rejection risk, while stable financial conditions favour approval.

Interpretability is supported by the compact rule set and concentrated feature usage, allowing decisions to be explained through a small number of relevant rules. Unlike logistic regression, which combines features linearly, RASF-LCS provides explicit, human-readable reasoning, making it suitable for regulated financial applications.

6 Conclusion

This paper presented Ranked Attribute Selection with Midpoint Filtering, an enhancement to Learning Classifier Systems (LCS) for credit approval tasks. By integrating feature ranking, rank-guided attribute selection, and distance-based filtering, RASF improves the focus and quality of learned rules. Experimental results show that RASF-LCS achieves a 3–5% accuracy improvement over standard LCS, which is meaningful in credit risk contexts.

Although Logistic Regression attains slightly higher accuracy, RASF-LCS provides a key advantage in interpretability through explicit, human-readable rules. This transparency enables auditing and supports decision accountability, which is essential in regulated financial environments. The learned rules also reveal intuitive patterns, e.g. the importance of default history and income-related features.

Limitations include sensitivity to hyperparameters and higher training cost compared to simpler models. Future work will focus on rule compaction, scalability to larger real-world datasets, and incorporating fairness-aware mechanisms.

7 Acknowledgments

The authors used ChatGPT to assist in improving the writing quality of a few sections. The authors remain fully responsible for the originality, accuracy, and overall integrity of the work.

References

[1] Hussain Abdou and Jonathan Pointon. 2011. Credit Scoring, Statistical Techniques and Evaluation Criteria: A Review of the Literature. *Intelligent Systems in Accounting, Finance and Management* 18, 2-3 (2011), 59–88. doi:10.1002/isaf.325

[2] Isidro M. Alvarez, Trung B. Nguyen, Will N. Browne, and Mengjie Zhang. 2025. A Layered Learning Approach to Scaling in Learning Classifier Systems for Boolean Problems. *Evolutionary Computation* 33, 1 (2025), 115–140. doi:10.1162/evco_a_00351

[3] E. Angelini, G. di Tollo, and A. Roli. 2008. A Neural Network Approach for Credit Risk Evaluation. *The Quarterly Review of Economics and Finance* 48, 4 (2008), 733–755. doi:10.1016/j.qref.2007.04.001

[4] Jason T. Bishop, Matthew Gallagher, and W. N. Browne. 2022. Pittsburgh Learning Classifier Systems for Explainable Reinforcement Learning: Comparing with XCS. *Proc. Genetic and Evolutionary Computation Conference* (2022), 323–331. doi:10.1145/3512290.3528767

[5] Matthias Bücker, Gero Szepannek, Alicja Gosiewska, and Przemysław Biecek. 2022. Transparency, Auditability, and Explainability of Machine Learning Models in Credit Scoring. *Journal of the Operational Research Society* 73, 1 (2022), 70–90. doi:10.1080/01605682.2021.1922098

[6] Michael Bücker, Gero Szepannek, Alicja Gosiewska, and Przemyslaw Biecek. 2022. Transparency, auditability, and explainability of machine learning models in credit scoring. *Journal of the Operational Research Society* 73, 1 (2022), 70–90. doi:10.1080/01605682.2021.1922098

[7] Thomas M. Cover and Joy A. Thomas. 1991. *Elements of Information Theory.* Wiley, New York. First edition.

[8] Xolisa Dastile, Tansel Celik, and Murendeni Potsane. 2020. Statistical and Machine Learning Models in Credit Scoring: A Systematic Literature Survey. *Applied Soft Computing* 91 (2020), 106263. doi:10.1016/j.asoc.2020.106263

[9] Maria Dumitrescu, Salvador Hué, Christophe Hurlin, and Sedley Tokpavi. 2022. Machine Learning for Credit Scoring: Improving Logistic Regression with Non-Linear Decision-Tree Effects. *European Journal of Operational Research* 297, 3 (2022), 1178–1192. doi:10.1016/j.ejor.2021.06.053

[10] Haifeng Gao, Gang Kou, Hu Liang, Hong Zhang, Xueli Chao, Chak-chun Li, and Yingfan Dong. 2024. Machine Learning in Business and Finance: A Literature Review and Research Opportunities. *Financial Innovation* 10, 1 (2024), 86. doi:10.1186/s40854-024-00629-z

[11] Muhammad Iqbal, Will N. Browne, and Mengjie Zhang. 2014. Reusing Building Blocks of Extracted Knowledge to Solve Complex, Large-Scale Boolean Problems. *IEEE Transactions on Evolutionary Computation* 18, 4 (2014), 465–480.

[12] Trung B. Nguyen, Will N. Browne, and Mengjie Zhang. 2023. ConCS: A Continual Classifier System for Continual Learning of Multiple Boolean Problems. *IEEE Transactions on Evolutionary Computation* 27, 4 (2023), 1057–1071. doi:10.1109/TEVC.2022.3210872

[13] Cynthia Rudin. 2019. Stop Explaining Black Box Machine Learning Models for High-Stakes Decisions and Use Interpretable Models Instead. *Nature Machine Intelligence* 1, 5 (2019), 206–215.

[14] Abubakar Siddique and Will Browne. 2022. Learning classifier systems: cognitive inspired machine learning for eXplainable AI. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion.* 1081–1110.

[15] Abubakar Siddique and Will N. Browne. 2022. Learning classifier systems: cognitive inspired machine learning for eXplainable AI. *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (2022). https://api.semanticscholar.org/CorpusID:250645203

[16] Abubakar Siddique, Will N Browne, and Gina M Grimshaw. 2020. Lateralized learning for robustness against adversarial attacks in a visual classification system. In *Proc. Gen. and Evol. Comput. Conf.* 395–403.

[17] Abubakar Siddique, Will N Browne, and Gina M Grimshaw. 2021. Frames-of-reference-based learning: Overcoming perceptual aliasing in multistep decision-making tasks. *IEEE Transactions on Evolutionary Computation* 26, 1 (2021), 174–187.

[18] Abubakar Siddique, Will N. Browne, and Gina M. Grimshaw. 2023. Lateralized Learning to Solve Complex Boolean Problems. *IEEE Transactions on Cybernetics* 53, 11 (2023), 6761–6775. doi:10.1109/TCYB.2022.3166119

[19] Patrick O Stalph and Martin V Butz. 2012. Guided evolution in XCSF. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation.* 911–918.

[20] Ryan J. Urbanowicz and William N. Browne. 2017. *Introduction to Learning Classifier Systems.* Springer.

[21] R. F. Zhang and R. J. Urbanowicz. 2020. A Scikit-Learn Compatible Learning Classifier System. In *Proc. Genetic and Evolutionary Computation Conference Companion.* 1816–1823. doi:10.1145/3377929.3398097