GPU-Accelerated Rule Evaluation and Evolution

Hormoz Shahrzad hshahrzad@cs.utexas.edu The University of Texas at Austin; Cognizant AI Labs, San Francisco, California

ABSTRACT

This paper presents an extension to the EVOTER (Evolution of Transparent Explainable Rule Sets) framework, previously introduced to promote transparency and explainability in AI through the evolution of interpretable rule sets. While EVOTER demonstrated strong performance in generating compact and understandable rulebased models using a list-based grammar, the current work, Accelerated Evolutionary Rule-based Learning (AERL), builds upon this foundation by addressing scalability challenges. Specifically, AERL leverages GPU-accelerated computation and back-propagation finetuning within a PyTorch framework to significantly enhance the efficiency of rule evaluation and evolution. By introducing a tensorized numerical representation of rules and incorporating gradient-based optimization, AERL achieves superior computational speed and improved search-space exploration. Experimental results confirm that the approach is effective, achieving faster convergence and similar accuracy. It thus offers a scalable path forward for explainable AI.

ACM Reference Format:

Hormoz Shahrzad and Risto Miikkulainen. 2025. GPU-Accelerated Rule Evaluation and Evolution. In *Genetic and Evolutionary Computation Conference (GECCO '25 Companion), July 14–18, 2025, Malaga, Spain.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3712255.3736397

1 INTRODUCTION

Evolutionary Rule-based Machine Learning (ERL) systems are extensively used in supervised learning, where rule fitness is typically measured by classification accuracy on labeled datasets [10]. A recent contribution in this area is EVOTER, which evolves expressive rule sets using a list-based grammar [11].

This work extends EVOTER by addressing a key bottleneck in ERL: the high cost of rule evaluation. GPUs have long accelerated optimization methods—including Mixed Integer Non-Linear Programming (MINLP), Central Force Optimization (CFO), Genetic Algorithms (GA), and Particle Swarm Optimization (PSO)—achieving speedups from $7 \times$ to $10,000 \times [13]$, with promising results reported across domains [8]. Building on this potential, this paper proposes a tensor-based rule representation in PyTorch, enabling parallel GPU evaluation and gradient-based fine-tuning via back-propagation for rule-based systems.

GECCO '25 Companion, July 14-18, 2025, Malaga, Spain

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-1464-1/2025/07 https://doi.org/10.1145/3712255.3736397 Risto Miikkulainen risto@cs.utexas.edu The University of Texas at Austin; Cognizant AI Labs, San Francisco, California

 Features (only six included for simplicity):

 mean radius, fractal dimension error, worst compactness,

 worst symmetry, perimeter error, mean area

 Rule 1: (Rule 1 has three conditions)

 Conditions: tensor([[10, 0, 0, 0, 0, 0], [0, 1.0, -0.38, 0, 0, 0], [0, 0, 0, 1.0, 0, 0]]),

 Biases: tensor([-2.02, 0.0, -0.29]), # 0.0 shows feature-to-feature comparison

 Rule 2: (Rule 2 has two conditions)

 Conditions: tensor([[0, 0, 0, 0, 1.0, 0], [0, 0, 0, 0, 0, 1.0]])

 Biases: tensor([-0.43, -0.29])

(a) Numerical representation of the rule set before fine-tuning

. (mean radius > 2.02) &
(fractal dimension error > $0.38 \times$ worst compactness) &
(worst symmetry > 0.29) \rightarrow malignant,

2. (perimeter error > 0.43) & (mean area > 0.29) \rightarrow benign

(b) Rule set before fine-tuning

1. (mean radius > 0.06) &	
$(-fractal dimension error > -7275.88 \times worst compactness) \&$	
(worst symmetry > -3.69) \rightarrow malignant,	
2. $(-perimeter \ error > -0.98)$ &	
$(-mean \ area > -0.24) \longrightarrow benign$	

(c) The same rule set after fine-tuning

Figure 1: Example of rule set representation and fine-tuning. A randomly generated rule set (not evolved) is shown in both numerical (tensor) form and human-readable form to illustrate how feature weights and biases translate into logical conditions. Back-propagation is applied to tune the coefficients, demonstrating how fine-tuning alone can substantially improve accuracy while keeping the logical structure intact.

The proposed approach, Accelerated Evolutionary Rule-based Learning (AERL), improves scalability and enhances search efficiency. Experiments on the UCI Breast Cancer dataset demonstrate significant speed gains.

2 BACKGROUND

Evolving transparent, explainable rule sets using extended propositional logic is a promising approach in explainable AI (XAI), offering interpretability that black-box models lack [6, 11]. While effective for compact rule discovery, scalability remains a challenge, particularly during computationally intensive evaluation phases.

GPUs have been widely used to accelerate evolutionary machine learning (EML). One such area is Learning Classifier Systems (LCS [15] and especially XCS [9]). When individual rules are evaluated and evolved within a population through the Michigan-style approach, GPU acceleration has shown promise in this process [2, 14]. More generally, techniques have been developed for speeding up evaluations in general evolutionary learning on general-purpose GPUs [4]. Recent surveys and meta-analyses [5, 12] emphasize growing interest in efficient hybrid symbolic-subsymbolic systems as well.

Similar gains have been more difficult to achieve in the Pittsburghstyle approach, [3], where each individual represents a complete rule set and is evaluated holistically. The structural mismatch between symbolic rules and GPU-optimized numerical formats is more severe, making it difficult to parallelize the evaluations. This

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

paper introduces a tensorized rule representation in PyTorch, enabling efficient GPU-based evaluation and backpropagation-based fine-tuning, and offering a Pittsburgh-style alternative to populationlevel approaches.

3 METHOD

AERL transforms list-based propositional logic expressions [11] into a tensorized form amenable to linear algebra and optimization. Rules are represented as inequalities, allowing matrix multiplication to evaluate multiple rules in parallel (Figure 1*a*). Each row in a rule's tensor corresponds to a condition, with non-zero entries indicating feature weights and the associated bias representing the threshold. For example, a row

[0, 1.0, -0.38, 0, 0, 0]

with a bias of 0 encodes the inequality

fractal dimension error > $0.38 \times$ worst compactness. This mapping is shown in Figure 1*b*, which translates the tensor form into a human-readable rule set. For inference, an activation function clips negative values to zero (false) and rounds positives to one (true), enabling binary decisions from continuous outputs.

The second innovation in AERL is fine-tuning with backpropagation. Since evaluation already occurs on GPUs, it is a way to take advantage of them more fully. A differentiable activation function (e.g., sigmoid) provides gradients for back-propagation, making it possible to adjust rule coefficients without altering logical structure.

Evolution proceeds via tournament selection and offspring generation with randomized coefficients. Back-propagation refines fitness at each generation, guiding evolution through the Baldwin effect [1]. Evolution runs on CPUs, with GPU experiments loading candidates for accelerated evaluation.

Figures 1*b* and 1*c* illustrate a random example rule set (not evolved) for the UCI Breast Cancer dataset: its accuracy improved from 21.05% to 92.40% after 50 epochs of fine-tuning. This example highlights how coefficient optimization alone—without evolutionary search—can substantially boost accuracy while keeping the logical structure intact.

4 EXPERIMENTAL RESULTS

4.1 Experiment Setup

The experiments aim to demonstrate computational efficiency gains from the tensorized rule representation and GPU acceleration, using the UCI Breast Cancer dataset as a well-known and accessible benchmark [16]. As the focus is on acceleration rather than classification accuracy, this dataset suffices to illustrate the method's effect; the approach is compatible with EVOTER and systems using similar grammars, and can be readily scaled to larger datasets.

Experiments compared GPU vs. CPU performance with and without back-propagation-based fine-tuning. Each setup was run 10 times with a fixed pool of 100 candidate rule sets on Google Colab T4-GPU. Fine-tuning applied 50 epochs of gradient updates to rule coefficients.

4.2 Evolutionary Performance

The performance of the PyTorch rules was first measured in terms of accuracy improvements over generations. As shown in Figure 2,



Figure 2: Evolutionary Progress of PyTorch Rules in the UCI Breast Cancer Domain with vs. without Fine-tuning. Each point is averaged over 10 runs, with 95% confidence intervals. In both cases, rule structures are evolved, but in the fine-tuning setting, back-propagation is applied at each generation to optimize rule coefficients without altering structure. This gradient-based refinement improves accuracy and accelerates search via the Baldwin effect $(p = 2 \times 10^{-7})$.



Figure 3: Comparison of GPU to CPU Acceleration Ratios with and without Fine-tuning. Plots show average CPU/GPU wallclock time ratio over 10 runs with 95% confidence intervals. Both with and without fine-tuning, GPU acceleration provides a minimum fivefold speedup compared to CPUs. Observed variations are due to the total number of rules and conditions per generation within a fixed pool size of 100 candidates. These results demonstrate that GPUs significantly accelerate rule-set evolution with and without fine-tuning.

the application of fine-tuning significantly expedited the evolutionary progress, achieving statistically significant performance improvements ($p = 2 \times 10^{-7}$). This result highlights how the method leaves the search for logical structure to evolution while reducing the search space by fine-tuning rule coefficients to their best values within each structure.

4.3 Computational Efficiency

Figure 3 illustrates the acceleration ratios achieved through GPU computation. Notably, both with and without fine-tuning, the use of GPU resources consistently resulted in at least a fivefold speed increase over traditional CPU processing. This acceleration was observed despite variations in the number of rules and conditions per generation, highlighting the efficacy of GPU acceleration in expediting the evolutionary computation of PyTorch rules.

5 DISCUSSION AND FUTURE WORK

The proposed AERL framework offers a scalable approach to ruleset evolution by leveraging GPU acceleration without compromising model quality—results from CPU and GPU implementations are equivalent in terms of classification accuracy, with performance gains stemming solely from faster computation.

Future improvements include extending support for categorical features to broaden applicability across diverse domains. Incorporating advanced rule structures, such as time-lags and exponents, may capture richer temporal or nonlinear dependencies. Probabilistic outputs can be enabled via Softmax activation, offering finer-grained class predictions [7].

REFERENCES

- J. Mark Baldwin. 1896. A New Factor in Evolution. The American Naturalist 30, 354 (1896), 441–451. http://www.jstor.org/stable/2453130
- [2] Martin V. Butz, Pier Luca Lanzi, and Stewart W. Wilson. 2011. Fast prediction computation in learning classifier systems using CUDA. In Proceedings of the 13th annual conference on Genetic and evolutionary computation. ACM, 169–176. https://doi.org/10.1145/2001858.2001953
- [3] Alberto Cano, Amelia Zafra, and Sebastian Ventura. 2014. Parallel evaluation of Pittsburgh rule-based classifiers on GPUs. *Neurocomputing* 126 (2014), 45–57. https://doi.org/10.1016/j.neucom.2013.01.049
- [4] María A. Franco, Natalio Krasnogor, and Jaume Bacardit. 2010. Speeding up the evaluation of evolutionary learning systems using GPGPUs. In Proceedings of the 12th annual conference on Genetic and evolutionary computation. ACM, 1039–1046. https://doi.org/10.1145/1830483.1830672
- [5] Michael Heider, David Pätzel, Helena Stegherr, and Jörg Hähner. 2023. A Metaheuristic Perspective on Learning Classifier Systems. In *Metaheuristics for Machine Learning: New Advances and Tools.* Springer Nature Singapore, 73–98. https://doi.org/10.1007/978-981-19-3888-7_3
- [6] Erik Hemberg, Kalyan Veeramachaneni, Babak Hodjat, Prashan Wanigasekara, Hormoz Shahrzad, and Una-May O'Reilly. 2014. Learning Decision Lists with Lags for Physiological Time Series. In Third Workshop on Data Mining for Medicine and Healthcare, at the 14th SIAM International Conference on Data Mining. 82–87.
- [7] Babak Hodjat, Hormoz Shahrzad, Risto Miikkulainen, Lawrence Murray, and Chris Holmes. 2018. PRETSL: Distributed Probabilistic Rule Evolution for Time-Series Classification. Springer International Publishing, Cham, 139–148. https: //doi.org/10.1007/978-3-319-97088-2_9

- [8] Paul Jähne. 2016. Overview of the current state of research on parallelisation of evolutionary algorithms on graphic cards. In *GI-Jahrestagung*. https://api. semanticscholar.org/CorpusID:4866385
- [9] Tim Kovacs and Manfred Kerber. 2001. What makes a problem hard for XCS? In Advances in Learning Classifier Systems. Springer, 80–99. https://doi.org/10. 1007/3-540-44640-0_7
- [10] Risto Miikkulainen. 2023. Evolutionary Supervised Machine Learning. In Handbook of Evolutionary Machine Learning, W. Banzhaf, P. Machado, and M. Zhang (Eds.). Springer, New York.
- [11] Hormoz Shahrzad, Babak Hodjat, and Risto Miikkulainen. 2024. EVOTER: Evolution of Transparent Explainable Rule sets. ACM Trans. Evol. Learn. Optim. (Nov. 2024). https://doi.org/10.1145/3702651 Just Accepted.
- [12] Abubakar Siddique, Michael Heider, Muhammad Iqbal, and Hiroki Shiraishi. 2024. A Survey on Learning Classifier Systems from 2022 to 2024. In Proceedings of the Genetic and Evolutionary Computation Conference Companion. ACM, 1797–1806. https://doi.org/10.1145/3638530.3664165
- [13] Satvir Singh, Jaspreet Kaur, and Rashmi Sharan Sinha. 2014. A Comprehensive Survey on Various Evolutionary Algorithms on GPU. https://api.semanticscholar. org/CorpusID:36146830
- [14] Kean Ming Tan, Chuan-Kang Ting, and Kay Chen Tan. 2013. GPU-accelerated eXtended Classifier System. In 2013 IEEE Congress on Evolutionary Computation. IEEE, 1–8. https://doi.org/10.1109/CEC.2013.6557250
- [15] Ryan J. Urbanowicz and Jason H. Moore. 2009. Learning classifier systems: A complete introduction, review, and roadmap. *Journal of Artificial Evolution and Applications* 2009 (2009), 1–25. https://doi.org/10.1155/2009/736398
- [16] Matjaz Zwitter and Milan Soklic. 1988. Breast Cancer. UCI Machine Learning Repository. DOI: https://doi.org/10.24432/C51P4M.