Recent advances in evolutionary rule-based machine learning (2024 to 2026)

Connor Schönberner, Hiroki Shiraishi, Fumito Uwano, Michael Heider

Angaben zur Veröffentlichung / Publication details:

Schönberner, Connor, Hiroki Shiraishi, Fumito Uwano, and Michael Heider. 2026. “Recent advances in evolutionary rule-based machine learning (2024 to 2026).” In *GECCO ’26 companion: proceedings of the Genetic and Evolutionary Computation Conference Companion, July 13-17, 2026, San Jose, Costa Rica*, edited by Leonardo Trujillo and Ting Hu, 1201–11. New York, NY: ACM. https://doi.org/10.1145/3795101.3814713.

Recent Advances in Evolutionary Rule-Based Machine Learning (2024 to 2026)

Connor Schönberner
cos@informatik.uni-kiel.de
Kiel University
Kiel, Germany

Hiroki Shiraishi
hiroki.shiraishi.we@hitachi.com
Hitachi, Ltd.
Tokyo, Japan

Fumito Uwano
uwano@go.tuat.ac.jp
Tokyo University of Agriculture and Technology
Tokyo, Japan

Michael Heider
michael.heider@uni-a.de
Universität Augsburg
Augsburg, Germany

Abstract

Evolutionary Rule-Based Machine Learning (ERL) employs evolutionary algorithms to automatically identify, evolve, and learn rules that decompose complex problems into subproblems, which are then solved using rule-based machine learning. Those rules subdivide the observation space into niches in a human-comprehensible manner by identifying innate data patterns and thus regularly serve as an explainable AI (XAI) component in safety-critical systems, garnering stakeholder trust and enabling enhanced verification. The International Workshop on Evolutionary Rule-based Machine Learning (IWERL), formerly IWLCS, is one of the oldest continuously running workshops at GECCO, having its 29th edition in 2026, and functions as a central forum for research and collaboration in ERL, serving both the next generation of researchers and the community as a whole. Following the tradition of previous surveys at the workshop, this work provides a systematic literature review of ERL-related publications from all sources in the years 2024 to 2026. We categorize each publication according to its background, methodology, and application into one of seven categories and summarise them. In doing so, this survey provides an accessible overview of recent advances in ERL, benefiting future work in this field and identifying research gaps to propose where the community could head in the future.

CCS Concepts

• **Computing methodologies** → **Rule learning**; **Genetic algorithms**; • **General and reference** → *Surveys and overviews.*

Keywords

Rule-based Learning, Learning Classifier Systems, Evolutionary Computation, Machine Learning, eXplainable AI, Survey

ACM Reference Format:
Connor Schönberner, Hiroki Shiraishi, Fumito Uwano, and Michael Heider. 2026. Recent Advances in Evolutionary Rule-Based Machine Learning (2024 to 2026). In *Genetic and Evolutionary Computation Conference (GECCO Companion '26), July 13–17, 2026, San Jose, Costa Rica.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3795101.3814713

ACM ISBN 979-8-4007-2488-6/2026/07
https://doi.org/10.1145/3795101.3814713

1 Introduction

The family of Evolutionary Rule-Based Machine Learning (ERL) methods employs evolutionary algorithms to automatically identify, evolve, and learn rules that decompose complex problems into subproblems, which are then solved using rule-based machine learning. These rules collectively provide knowledge to solve the learning problem in an ensemble-like fashion. Evolutionary Machine Learning is continuing to grow in popularity due to its combination of open-ended population-based problem solving techniques with effective machine learning. ERL [22, 77] is combining these concepts to identify innate patterns in the data it encounters in a human-comprehensible manner and thus can serve as a key component towards explainable AI (XAI) systems [24, 90].

The International Workshop on Evolutionary Rule-based Machine Learning (IWERL), formerly known as the IWLCS, is one of the if not the oldest running workshop at GECCO, having its 29th edition in 2026. It plays an important role as a forum for research and collaboration in evolutionary rule-based machine learning, with a traditionally strong influence of Learning Classifier Systems (LCSs), serving both the next generation of researchers and the community as a whole. Since 2019, organizers of IWERL have been providing a comprehensive overview of ERL research that was recently conducted [e.g. 23, 43, 47, 68]. The primary goal of this work is to provide a comprehensive survey on using ERL techniques over the past two years. By doing that, we intend to contribute to a better organized research community; showcasing the most recent developments of the field helps to connect ERL researchers but also serves people new to the field as they can more quickly assess the latest achievements, current challenges as well as links to other areas.

In the next section, we shortly explain ERL and follow with a section describing the methodology we used for identifying publications to be included in our overview. All the 63 selected papers are divided into major categories based on the nature of their contributions. The following seven sections correspond to the aforementioned categories, one section on identified research tendencies and research gaps and a concluding summary. We mention each publication only once, even if some fit into several sections; for each paper we choose the section it fits into best. If an equally strong case for multiple sections could be made, we choose the first appropriate section, in the order in which the sections appear.

2 ERL in General

ERL encompasses several prominent categories, such as LCSs, Ant-Miner, Artificial Immune Systems, and fuzzy rule-based systems. The modes or model structures of these systems are optimized using evolutionary, symbolic, or swarm-based methods and the basic structure of the models are lists or sets of rules that approximate a complex decision landscape by dividing it into smaller subproblems. They have been successful in all major machine learning paradigms.

Similar to how Drugowitsch [14] formalised LCSs as generalised probabilistic mixture of expert models, Heider et al. [22] generalised the semantics of ERL models to one single formalisation encompassing most to-date known ERL architectures, regardless of training approach or application domain. Predictions of an ERL method's global model, denoted as $\hat{f}$, can be characterised as [19, p. 16]:

$$\hat{f}_{\mathcal{M},\theta}(\mathbf{x}) = \sum_{k=1}^{K} m(\psi_k; \mathbf{x})\, \gamma_k\, \hat{f}_k(\theta_k; \mathbf{x}) \tag{1}$$

where

- $\mathbf{x}$ is the input for which a prediction is to be made,
- K is the overall number of rules in the model, which, together with the set of matching function parameters $\psi_k{}_{k=1}^{K}$, forms the *model structure* $\mathcal{M} = (K, \{\psi_k\}_{k=1}^{K})$,
- $m(\psi_k; \mathbf{x})$ is the output of the matching function of rule k defined by its parameters ψ_k for an arbitrary input $\mathbf{x}$; typically, $m(\psi_k; \mathbf{x}) = 1$ iff rule k is responsible for input $\mathbf{x}$ (i.e. if its local model models that input), and $m_{\psi_k;k}(\mathbf{x}) = 0$ otherwise [note: in Fuzzy-ERL, this is generalized to a continuous membership degree in $[0, 1]$; cf. Section 6],
- γ_k is the mixing weight of rule k; mixing weights usually fulfil $\sum_{k=1}^{K} m_k(\mathbf{x})\gamma_k = 1$ and $0 \le \gamma_k \le 1$ (this can easily be achieved using normalization, which is standard for almost all LCSs),
- $\hat{f}_k(\theta_k; \mathbf{x})$ is the output of the *local model* (sometimes also referred to as a *submodel*)—a discriminative function $\hat{f}_k(\theta_k) : \mathcal{X} \rightarrow \mathcal{Y}$—of rule k (with parameters θ_k) for input $\mathbf{x}$,
- and $\theta = \{\theta_k, \gamma_k\}_{k=1}^{K}$ forms the (whole) model's *parameters* out of the local models' individual parameters θ_k and their corresponding mixing weights γ_k.

3 Methodology

For identifying works for this survey, we limited our systematic searching to contributions in English with *publication dates* on or after 11 March 2024 (one day after the end of the period of the previous survey) and on or before 03 February 2026. Note that some conferences may have published their proceedings within this interval despite in fact taking place at another date; since we only consider publication dates, these entries are covered by this survey as well. For this survey, we excluded Artificial Immune Systems to narrow the scope. In addition, papers not employing an evolutionary method in some form were excluded. Also, we leave out papers published on *arXiv* unless they are relevant to a follow-up paper which meets our criteria. Finally, we exclude publications that were already cited in the previous surveys even if they were eligible for the present survey as well (this may be the case, for example, if they have been re-published or moved from an 'accepted' to a 'published' status, thus stating a new date).

Our primary search tool was *Google Scholar*[1] with its time range feature as it seemed to consistently report more relevant publications than comparable tools, especially for more general search queries. We ultimately used seven different queries, three very general ones and one for each of the three major ERL systems that we knew had been investigated in recent years; more specific terms did not yield relevant results that were not found by these as well (at least not in February 2026). The list below shows each query with the number of Google Scholar result pages that we examined for papers meeting our requirements (numerator) and the total number of result pages (denominator); each Google Scholar result page displayed 10 publications and we set the time range filter to *Since 2024*.

- `learning "classifier system"` (34/>100)[2]
- `"evolutionary rule-based" learning` (6/6)
- `xcs classifier system` (31/31)
- `"exstracs" classifier system` (4/4)
- `antminer` (34/34)
- `ACS anticipatory "classifier system"` (2/2)

We also checked a number of ERL researchers (especially, but not exclusively, ones we know to be active) and the proceedings of the four main venues that have attracted ERL researchers in the past: CEC, GECCO, PPSN, and *EvoStar*. This led to no additional publications beyond those already identified on Google Scholar. Hence, we consider our methodology to be sufficiently exhaustive and did not refine the search terms.

4 Theory & Meta analysis

Theoretical advancements and unifying perspectives remain an object of research for ERL.

Schröder and Textor [55] revealed a deep theoretical connection between LCSs and probabilistic models by formalizing them as generalized linear probability models. This novel perspective mathematically demonstrates that an LCS can theoretically represent any arbitrary discrete probability distribution. Furthermore, this reframing allows for efficient conditioning and marginalization operations. Empirical evaluations on sequence modelling showed that simple LCSs exhibit a complexity-calibration trade-off comparable to Markov chains.

To address the computational bottleneck of evaluating massive populations, Schröder et al. [56] introduced a compression methodology using weighted finite state machines (WFSMs) with exact rational arithmetic. By representing weighted sets of strings compactly, this approach reduces memory usage and runtime by orders of magnitude. While successfully demonstrated on immune-inspired anomaly detectors and evolutionary algorithms (EAs), the authors explicitly emphasize that this WFSM framework is highly applicable for scaling up LCSs.

[1] https://scholar.google.com

[2] We stopped when we noticed that the results of at least ten consecutive pages did not contain any LCS-related publications that had not yet been listed on the preceding pages (i. e. we estimated the probability of further relevant results to be very low).

Addressing the perceived complexity of LCSs, Lanzi and Loiacono [29] proposed a simplified, modern narrative for the XCS Classifier System (XCS). They reframed XCS strictly as a rule-based action-value function approximator. Their approach describes the system using three basic functions: (i) XCS(θ), which initializes an instance using a parameter vector θ; (ii) ActionValues(s), which computes the prediction array for state s; and (iii) Update(s, a, P), which updates matching classifiers that advocate action a using the expected payoff P. This leaner (Q-learning-like) formulation facilitates easier understanding and seamless integration with contemporary reinforcement learning techniques.

5 New ERL architectures

Aside of extensions to existing systems, works regularly establish novel ERL architectures that differ substantially from prior systems.

The Heuristic Evolutionary Rule Optimization System (HEROS) introduced by Lipschutz-Villa et al. [30] is a two-phase LCS for supervised learning similar to SupRB [25]. Phase I corresponds to a Michigan-style LCS (descended from XCS/sUpervised Classifier System (UCS)/Extended Supervised Tracking and Classifying System (ExSTraCS 2.0)) with a novel subsumption and prediction array strategies, expert-guided covering and mutation, as well as a mixed discrete–continuous attribute representation similar to ALKR [5, 76]. Phase II works more similar to a Pittsburgh-style LCS employing NSGA-II [10] inspired multi-objective Pareto optimization to form rule-sets. It features a custom merge operator for its Genetic Algorithm (GA) and a model initialization based on the randomness and target accuracy trade-off. HEROS outperforms ExSTraCS 2.0 on the noisy GAMETES classification datasets using in some instances only 10% of the population size. However, ExSTraCS 2.0 performed better on three of five investigated multiplexer tasks, while HEROS still evolved vastly smaller populations consistent with its specialization for noisy data.

Liu et al. [33] introduced Absumption and Subsumption Learning Classifier System Real Plus (ASCS$_R^+$), an improved phenotypic LCS following their prototype [32]. Instead of evolving rule genotypes, its EA evolves phenotypes (covered instance sets), reducing overfitting via absumption and subsumption on heterogeneous instance sets and introducing informed mutation for more general rules. They show ASCS$_R^+$ substantially outperforms UCS and XCS for real input (XCSR) on 2048-feature datasets with higher accuracy and smaller populations, outperforming other white-box methods and rivalling black-box models, though still outperformed by large deep learning models.

Evolution of Transparent Explainable Rule-sets (EVOTER) by Shahrzad et al. [58] is a rule-set evolving system for supervised learning and policy-searching tasks such as Reinforcement Learning (RL) environments. Their rules consist of conditions and actions. EVOTER employs list-based propositional logic expressions as conditions that chain coefficients multiplied with features via operators. The rules' conditions, actions and coefficients are optimized by a GA. EVOTER either executes the action of the first matching rule, one determined by a hard-max filter on the action coefficients or all actions of the matching rules either sequentially or in parallel. AERL [59] extends EVOTER by introducing tensor-based rules powered by PyTorch for faster wall-clock training. It translates the list-based conditions to a tensor representation enabling more GPU-accelerated evaluation, while the GA operations are still fully CPU-bound. They pair this with gradient-based activation functions for fine-tuning the coefficients of their rules. Shahrzad and Miikkulainen [59] report higher accuracy in the UCI Breast Cancer classification task and improved wall-clock performance.

Liu et al. [31] propose Consistent Interval Decision Set (CIDS), a rule-based classification method for continuous-feature problems that is related to evolutionary rule-based machine learning. CIDS constructs decision sets directly from training data by first decomposing a multiclass task into one-vs-rest binary problems and then further dividing each binary problem into simpler sub-problems, which are converted into interval-based rules using the IMR operator and the TC-CNF rule representation. The authors argue that, compared with XCSR's stochastic rule search, CIDS yields more accurate and more interpretable rule sets on complex continuous benchmarks by using a divide-and-conquer strategy guided by data distribution.

Tzeng et al. [75] proposed the multi-agent RL framework Exanna incorporating values as motivational bases and rationales and apply their system to a pandemic simulation. Exanna trains a Q-learner to select actions incorporating the motivational values based on the beliefs derived from the environment and feedback from other agents. They do not specify whether it is a classic Q-learner or a DQN variant. In tandem, an XCS is used to generate rationales for the decisions of the Q-learner for updating the beliefs. Based on their results, the Tzeng et al. [75] conclude Exanna leads to higher social experience, conflict resolution, and flexibility.

Kovas and Hatzilygeroudis [27] provide an overview of their ACRES systems introducing two novel variants ACRES-SEQ and ACRES-4 as well as a new user interface. ACRES extracts rules from a dataset consisting of a condition and an estimated certainty factor to construct an expert system. Their two new approaches apply a GA to discover the feature pattern with the highest coverage and accuracy per class. From this pattern they construct a rule with a condition that captures the pattern with a corresponding sequence of feature value pairs. Taking uncertainty into consideration, each rule is augmented with a certainty factor, estimated with one of four proposed probability-based methods. The authors demonstrate that their ACRES variants reach comparable results to other examined Machine Learning (ML) methods.

6 Fuzzy-ERL

Fuzzy-ERL, particularly represented by Learning Fuzzy-Classifier Systems (LFCSs), combines the adaptive learning capabilities of EAs with the continuous logic of fuzzy systems. Unlike traditional crisp rule-based systems that divide data with strict boundaries (e.g., LCSs), fuzzy rules use continuous membership degrees, ranging from 0 to 1, which are defined by Membership Functions (MFs). This fuzziness allows the models to handle uncertainty, imprecise information, and noisy data much more effectively. Furthermore, by using linguistic terms (e.g., "small", "medium", or "large"), they can represent fuzzy rules that humans can intuitively understand.

To improve the decision-making process in LFCSs, Shiraishi et al. [63, 64] introduced a novel class inference scheme based on the Dempster-Shafer Theory of Evidence. The proposed scheme calculates "belief masses" for specific classes based on fuzzy rules. Importantly, it also includes an "I don't know" state, which explicitly quantifies uncertainty, unlike conventional inference schemes, such as weighted voting and single-winner. By combining these beliefs, the system provides transparent and reliable predictions, forms smoother decision boundaries, and significantly improves generalization performance on real-world datasets.

A highly flexible rule representation that automatically adapts to the complexity of different input subspaces was proposed by Shiraishi et al. [61, 62]. This representation, based on a four-parameter beta distribution (FBR), was integrated into a system called β_4-UCS. FBR enables MFs to change shape dynamically, ranging from crisp rectangles to fuzzy bell curves, depending on the specific needs of each local subspace. In this way, β_4-UCS can use simple, interpretable crisp rules where the decision boundaries are clear and complex, accurate fuzzy rules where they are not. This results in highly accurate, compact rulesets.

In addition, Shiraishi et al. [65] introduced a stochastic local mutation operator for Adaptive Fuzzy-Classifier System (AFCS) [60] to efficiently explore various types of MFs. While allowing AFCS to automatically choose from multiple MF types (like rectangular, triangular, or bell-shaped) increases model flexibility, it also massively expands the search space, which can slow down the learning process and degrade efficiency. The proposed mutation operator addresses this challenge by modelling the probability distribution of successful MF types extracted from promising rules. It probabilistically applies this local knowledge to create new rules, balancing the exploitation of good MF patterns with the exploration of new ones. This speeds up learning and improves final classification accuracy.

To handle higher-dimensional data more effectively, Shiraishi et al. [67] proposed Fuzzy-UCSv with a variable-length fuzzy set representation. Traditional LFCSs (e.g., Fuzzy-UCS) require exactly one MF for every feature, which wastes computational effort on irrelevant features and limits a single rule to covering only one subspace. Fuzzy-UCSv solves this limitation by allowing the number of MFs per feature to vary. This means the system can completely ignore non-contributing features (acting as "Don't Care") and use multiple MFs to cover complex, disconnected subspaces with just a single rule. This variable-length representation significantly improves testing accuracy and prevents overfitting, especially when the maximum population size is severely restricted.

7 ERL and neural networks

Deep learning offer impressive computational as well as representational learning capabilities but remain essentially black-boxes.In contrast, ERL methods offer interpretability of their knowledge base and the decision-making process. Hybrid methods promise, in the best case, to combine the benefits of both worlds.

Previous work established neural networks replacing the condition and action, termed as neural rules [7], or only the condition, termed as neural conditions [50]: Schönberner et al. [52] examine modernised versions of both concepts in multi-step RL environments. They find that neural rules with binary encoded actions and neural conditions can improve over hyperrectangles when employed for XCS for function approximation (XCSF) in multi-step RL environments.

Tianrungroj and Iba [73] introduced a Large Language Model (LLM) and LCS hybrid coined Rule Optimization method using LLMs for Imitation Learning (ROIL). As LCS, they employ their novel Natural Language-based LCS. They improved over their chosen baseline ReAct [85] – a well-regarded LLM approach for text-based interactive tasks. Within ROIL, the LLM is employed in the matching logic and as the optimiser, where gradient-free prompt fine-tuning is used for Imitation Learning. Base ROIL employs no conventional evolutionary optimisation. An additional variant iteratively improves the formed rule set with a mutation-only EA, but the authors report that this so far led to no improvement.

Yatsu et al. [87] addressed the curse of dimensionality in high-dimensional prototype generation by combining two of their previously proposed methods: a prototype generation algorithm named kNNUCS-PG [88] (cf. Section 8.2) and a latent space exploration strategy named LS-OvL [86]. By integrating these with a Variational Autoencoder (VAE), this approach evolves rule representations in a low-dimensional latent space and evaluates decoded prototypes in the original high-dimensional observation space. This synergy significantly reduces the number of generated prototypes while achieving higher classification accuracy than existing methods.

Focusing on dimensionality reduction for visual RL on image-based grey-scale states, Schönberner et al. [51] combined XCSF with a VAE to the AEXCSF. Firstly, they trained the VAE offline on states collected by another RL method, then XCSF was trained online on the latent representations of that VAE (mean μ, μ and σ concatenated, or re-parametrised states). As a simple baseline they downscaled images. For the VAE approach μ matched or beat concatenated μ and σ while reparametrisation underperformed. Downscaling proved to be a simple yet effective alternative, but the VAE approach learned smaller representations than the downscaled image size. As a result, they projected that the VAE approach should outperform downscaling in sufficiently complex environments where downscaling loses too much information.

Shiraishi et al. [66] introduced a function approximation method called X-KAN, which integrates Kolmogorov-Arnold Networks (KANs) into the evolutionary rule-based framework of XCSF. While standard global Neural Networks (NNs) (e.g., Multi-Layer Perceptrons (MLPs), KANs) can struggle with locally complex and/or discontinuous functions, X-KAN overcomes this limitation by adaptively partitioning the input space and optimizing multiple localized KAN models as rule consequents. Specifically, X-KAN uses an EA to discover general local regions and backpropagation to train the local KAN. This divide-and-conquer approach significantly outperforms state-of-the-art methods, including the global MLP and KAN, in terms of approximation accuracy, all while maintaining a remarkably compact ruleset. Furthermore, since KANs are known to function as symbolic regressors, X-KAN is expected to be utilized as an interpretable piecewise symbolic regressor in the future.

8 Improvements to existing algorithms

Several works of the past two years concentrated on improving or extending existing ERL systems or algorithms. Due to finding more than one contribution to each of them, we decided to give XCS with Section 8.1, UCS and ExSTraCS 2.0 with Section 8.2, Anticipatory Classifier System 2 (ACS2) with Section 8.3, Supervised Rule-based learning system (SupRB) with Section 8.4 and the Ant-Miner with Section 8.5 their own subsections. Further papers in this category are grouped in Section 8.6.

8.1 Improvements and extensions for XCS

XCS [82] is an accuracy-based, Michigan-style LCS applicable to RL, supervised learning and unsupervised learning domains. It combines a GA with a Q-learning-like RL mechanism. Following Wilson's generalisation hypothesis, XCS evolves maximally general classifiers. For real-valued inputs, XCS typically employs a hyperrectangular rule representation as a sequence of lower and upper bounds for each input. It is often also extended with computed linear prediction. XCS serves as a foundational architecture for many later systems.

Schönberner and Tomforde [54] investigated the impact of an activated or deactivated GA in multi-step RL experiments using XCSF with linear Recursive Linear Least Squares (RLS) prediction and neural prediction employing hyperrectangle conditions. Linear prediction needed the GA for sufficient performance, but neural prediction matched or exceeded it relying only on random covering populations (less than 60 rules on average). Since covering populations should be almost never optimal, they concluded that this indicates research potential for the GA, particularly in encouraging smaller rule populations.

Schönberner and Tomforde [53] introduce three Deep Q Network (DQN)-inspired [37] extensions for XCSF in multi-step RL: Double Q-learning [81] (reduced performance), target prediction [37] for stability, and double target prediction adapting Double DQN [80] to reduce overestimation. For linear prediction with RLS, double target or target prediction slightly to moderately improved performance in most targeted environments. This was mostly also the case with neural prediction. In particular when employing Experience Replay for XCS [70], results suggest positive impact of the extensions due to improved stability and better performance.

Uwano and Browne proposed Aliasing-Group-based FoRsXCS (AGFX) to identify ambiguous states that appear in regular patterns in multi-step decision-making tasks. AGFX forms aliasing groups from adjacent aliased states to capture such regularities, enabling the agent to identify the same type of ambiguous states across different aliasing groups [79]. Furthermore, they integrated AGFX with dual-stream identification (AGFX2) [78], which estimates the confidence of disambiguated states. AGFX2 helps prevent excessive rule deletion in multi-step decision-making tasks, thereby facilitating the formation of an optimal policy in complex aliasing tasks.

Following the dynamic-sized experience replay strategy introduced in LCT-DER (Learning Classifier Table with Dynamic-sized Experience Replay), Surhonne et al. [72] proposed XCS-DER for memory-constrained XCS. XCS-DER determines the replay memory size based on the current macro-classifier population, thereby reallocating unused memory from the classifier population to stored experiences under a fixed memory budget. Their results on single-step benchmark tasks suggested better learning efficiency than both standard XCS and XCS with a fixed-size replay memory.

Siddique et al. [69] propose AXCSR, an attention-oriented extension of XCSR for visual classification, in which an input image is first split into a fixed set of patches and features are extracted from each patch, after which the system learns rules that determine which patch-level information should be attended to and which should be ignored through don't-care encoding. As a proof of concept, the method is evaluated on cat-versus-dog image classification using LBP, HOG, GLCM, SIFT, and their concatenation. Although AXCSR generally learns more slowly than standard XCSR in higher-dimensional settings, it still achieves comparable final accuracy in several cases, while GLCM alone remains weak for both methods.

Novak et al. [40] proposed an exploratory framework that combines numerical association rule mining with XCS by injecting pre-mined rules into the initial classifier population. In this approach, numerical attributes are discretized and binary-encoded so that both dataset instances and mined rules can be handled by a basic XCS. Their results suggest that such rule injection can improve performance under a simplified reward setting, but the formulation is rather unconventional because dataset attributes are treated as actions, making the connection to standard XCS indirect.

Alvarez et al. [3] extend XCSCFA for Boolean domains by combining layered learning and transfer learning, where each target problem is manually decomposed into a sequence of subordinate tasks and the knowledge acquired at earlier stages is reused in later ones through code fragments and learned rule-set functions. Their framework introduces typed code fragments, an attribute-list primitive, and constant terminals to make the reuse mechanism more robust across domains, and it is evaluated on Multiplexer, Carry-one, Majority-on, and Even-parity problems.

Pour-Hosseini et al. [49] present TinyXCS, a lightweight XCS-based controller for workload allocation in edge computing under resource constraints. The method models the current workload, network congestion, and battery level as the decision context, and uses XCS to select the amount of workload to be processed locally at the edge, with the remainder offloaded upstream. Technically, the paper mainly instantiates a standard XCS pipeline for this resource-management setting, while emphasizing reduced memory usage rather than a new classifier representation or learning mechanism.

8.2 Improvements and extensions for UCS and ExSTraCS

UCS is an accuracy-based, Michigan-style LCS designed specifically for supervised classification tasks and is a simplified version of XCS. By combining a GA with supervised learning mechanisms, UCS aims to evolve a minimal yet highly accurate and maximally general ruleset online. For real-valued inputs, UCS typically employs a hyperrectangular rule representation defined by lower and upper bounds for each feature. Because of its robust generalization capabilities and interpretability, UCS serves as a foundational architecture for various modern extensions.

Focusing on data reduction while preserving classification accuracy, Yatsu et al. [88] proposed kNNUCS-PG, a novel prototype generation method that directly extends UCS. Departing from traditional hyperrectangular conditions, it employs a point representation alongside a k-nearest neighbour matching algorithm. By utilizing contribution-based and merge deletion mechanisms, the system effectively evolves representative prototypes near decision boundaries. This approach significantly reduces dataset sizes while maintaining high performance in artificial and real-world tasks.

In the domain of digital forensics, Todd and Peterson [74] introduced Digital Trace Inspector (DTI), a decision support tool using ExSTraCS—an advanced extension of UCS—to automate temporal metadata analysis. DTI utilizes a YARA-inspired expert knowledge framework to encode complex, interleaved metadata traces into interval and characteristic features. By evolving rulesets from these abstractions, DTI accurately identifies corroborating evidence, significantly reducing the time and effort required compared to manual forensic methods.

Woodward et al. [83] introduced Survival-LCS, a descendant system of an enhanced ExSTraCS 2.0, adapting their novel system to model time-to-event survival data. They further deepened their research in follow-up work [84] with broader hyperparameter search and a variety of survival distributions to assess the capabilities and limitations of this novel LCS. For their use-case, Woodward et al. [84] designed Survival-LCS for quantitative prediction by using batch-evaluation for newly evolved rules and model actions as intervals of form $[min\_time, max\_time]$. Overfitting is avoided and generalization improved, by incorporating a multi-objective rule fitness function to either optimize accuracy and/or correct instance coverage. Last but not least, they adapt the correct and incorrect set formation to cope with censored data which is crucial for the biomedical domain. They found that their Survival-LCS can learn complex associations with survival outcomes and performs competitively compared to standard methods, in particular in areas challenging the latter.

8.3 Improvements and extensions for ACS2

ACS2 [8] is an accuracy-based, Michigan-style Anticipatory LCS designed for multi-step RL. It combines a GA with Q-learning-like RL mechanism, while also integrating an anticipation component that predicts the follow-up state. Extended variants of ACS2 are in particular designed to deal with perceptual aliasing which is a well-known challenge for RL methods.

Orhand and Ayadi [41] introduce a de-biasing feature to ACS2 variants including BEACS, enabling first-time cart-pole solution for ALCSs when combined with discretizing states via bucketing, while maintaining performance across 23 maze maps. Furthermore, they introduce a neurosymbolic-ALCS framework coupling ACS2 variants with ontology-based supervision as semantic safeguards. Their two illustrated examples for maze maps demonstrate that ontologies can validate actions, rejecting dangerous ones, and can bias towards exploring low-experience states.

Łabędzki and Unold [28] show in their empirical benchmark study that despite advances in ACS2 variants (ACS2, PEPACS, BACS, BEACS) across 26 maze environments, perceptual aliasing persists, especially in complex mazes where even BEACS fails to fully distinguish aliased states. They suggest that a memory mechanism could address the differentiation issue and suggest that improvements of the generalization and subsumption mechanisms are needed for further progress.

Hayashida et al. [18] apply ACS2 in a multi-agent RL environment for pedestrian simulation, focusing on bottlenecks created by obstacles and corridors. Compared to Q-learning, ACS2 matches the learning performance while it provides superior stability and interpretability, enabling the analysis of collective decision-making and movement patterns.

8.4 Improvements and extensions for SupRB

The SupRB [25] is a multi-solution batch learning approach for training interpretable regression models. Its key advantage is that it directly optimizes the rule set complexity as well as its prediction error and can therefore construct substantially more compact models. This is achieved by separating rule discovery and model construction (called solution composition; SC) into independent optimization problems that occur in alternating phases and can therefore exploit each others' knowledge.

In [21], SupRB was further improved by replacing the SC phase's rather simple genetic algorithm with self-adaptive versions. In the experiments, four different adaptive optimization approaches were compared against the baseline. The key result was that these adaptive versions manage to produce even less complex yet more predictive models while cutting down on the number of hyperparameters that need tuning.

In his doctoral thesis [19], Heider summarized his previous work on SupRB and ERL in general. Additionally, a more in-depth perspective on real-world usage was given, as well as extended experiments on all major components including fully novel experiments regarding the mixing model (cf. Section 2) used in SupRB.

8.5 Improvements and extensions for the Ant-Miner

The Ant-Miner [42] employs Ant Colony Optimisation (ACO) to construct accurate rules for classification tasks.

In [38], some slight tweaks to the well-known Ant-Miner rule-based learning system are proposed and evaluated on a publicly available data quality dataset collected by a group from Australia. The authors show that their version uses slightly fewer rules while achieving a comparable accuracy.

Ayub and Almas [4] provide a more substantial improvement to the Ant-Miner with an improved ant selection and quality function that judges ants based on coverage, rule length, and accuracy. They show that their method outperforms or is competitive with several state-of-the-art Ant-Miners on four benchmark classification datasets in terms of model complexity and predictive power.

Brookhouse et al. [6] introduced Ant-Miner$_{\text{PB+HMA}}$, a hybrid-pheromone ACO classifier that removes dynamic discretization and replaces multi-pass pruning with a single-pass method, reducing runtime while maintaining accuracy. Building on this, it forms ensembles via feature bagging and bootstrapped colonies. Across 21 UCI datasets, they achieved competitive predictive performance and significantly faster execution on larger datasets than cAnt-Miner$_{\text{PB}}$, therefore, removing the limitations on dataset size that previous Ant-Miners faced.

8.6 Extending Other ERL Approaches and General Improvements

Pätzel et al. [44, 45] investigate variable-length metaheuristic rule-set learning for supervised regression. Each candidate solution is a rule-set composed of interval-based conditions and simple local models. The genetic algorithm optimizes the rule-set structure, such as the number of rules and their condition boundaries, while local rule parameters are fitted from data and the whole model is evaluated with an AICc-based fitness that balances predictive fit and compactness. In [45], an initial benchmark of spatial crossover and length-niching selection against more standard operators is presented. It shows that all variants achieve broadly similar test MAE on 54 synthetic regression tasks, but that avoiding crossover tends to produce more compact solutions. In [44], an extended follow-up analyses convergence behaviour in more detail, examines why crossover can hurt compactness, and adds comparisons with XCSF and CART. It concludes that one crossover variant may trade a few large errors for many smaller ones at the cost of larger rule sets, that the benefits of length-niching selection are not clearly visible in fitness on the tested tasks, and that the best GA variant outperforms XCSF and comes close to CART while maintaining similarly compact models.

Zang et al. [89] improved ZCSAR, a ZCS variant based on average-reward RL, by revising its selection mechanisms with minimal changes to the original framework. To address the difficulty that negative classifier strengths make standard roulette-wheel selection unreliable, they introduced truncation and compared several combinations of tournament and roulette-wheel selection for GA parent selection, action selection, and classifier deletion. Experiments on the Maze6 and Woods14 multistep maze tasks showed that using tournament selection only for GA parent selection, while retaining roulette-wheel selection for action selection and deletion, yielded the most stable learning behavior and the best or near-optimal performance across parameter settings.

9 Benchmarking

One central question when benchmarking ERL algorithms is the similarity of models. Commonly, this is assessed through measurements related to the model's complexity and aggregated prediction performance. For example, when comparing models, e.g. generated with different metaheuristic operators (e.g. mutation strategies), their mean squared error and the number of rules are used. This tells us whether models are similar with regards to metrics but not whether they "look" the same, i.e. whether their rules use similar conditions. Pätzel et al. [46] propose and successfully demonstrate the application of a new metric to mitigate this gap.

XCS has been a go-to method for constructing systems that respond autonomously and robustly to environmental changes. However, XCS is rarely benchmarked against other RL approaches. Steinberger et al. [71] present an extensive empirical analysis of XCS on 51 problem instances from five established RL benchmarks. The results show that while XCS performs competitively on simpler problems and is robust in noisy environments, it fails completely on two of the benchmarks.

Coombe et al. [9] evaluate LCSs on the abstraction and reasoning corpus, a challenging AI benchmark where unseen relationships ought to learn from few data points. They show that even simplistic LCSs (eLCS and XCS) can solve a variety of puzzles with kernels of 3x3 to 5x5 input size but fail to scale to larger kernels.

10 Applications

Over the surveyed period, there has been a variety of different successful applications of ERL methods.

Heider [20] argues, on the basis of an application in manufacturing, that when explainability is needed for an AI-based system, there still likely are differences in what will constitute a good model and specific requirements engineering is needed. It posits that SupRB (cf. Section 8.4) may be an ideal approach to fulfill the requirements in supervised learning settings based on the current research works on this system.

Motivated by an increasing complexity in cyber-physical networks, Feist et al. [16] present *Chameleon*, an adaptive middleware that operationalizes an LCS-driven MAPE-K loop to manage system- wide real-time constraints. They demonstrate that LCSs can be effectively deployed in resource-constrained cyber-physical environments through a scalable rule-evaluation mechanism, normalized health metrics, and a domain-specific, human-readable rule language (Rango [15]) that provides an effective encoding for input vectors. The results highlight the practical viability of LCS-based adaptation within realistic automotive simulations.

In the medical domain, Hamouda et al. [17] proposed an explainable, heterogeneous ensemble learning approach for automated sleep stage classification using polysomnographic recordings. To address the trade-off between accuracy and interpretability in traditional machine learning models, their framework combines accuracy-based LCSs (UCS and XCS) with various base classifiers through voting and stacking mechanisms. The stacking model achieved an 89.2% accuracy while providing explicit, human-readable rules, offering sleep specialists a reliable and transparent diagnostic tool.

Addressing the self-optimization of soft real-time systems, Masood et al. [34] introduced a continual reinforcement learning-based expert system for virtual globes. Using XCSCFC, the system evolves and transfers adaptation rules to dynamically control viewer navigation speed and display resolution. This enables memory-efficient knowledge reuse across single-incremental and multi-tasking scenarios, optimizing performance across diverse hardware capabilities without requiring manual pre-deployment tuning.

In the automotive domain, Nasim et al. [39] developed a cognitively inspired, sound-based automobile problem detection system to advance explainable AI. The system uses ExSTraCS to detect various engine faults by extracting time, frequency, and time-frequency domain features from engine acoustics. This rule-based approach outperforms conventional machine learning models by a significant margin, achieving 98.6% accuracy while generating transparent, human-interpretable rules for reliable vehicle diagnostics.

Akyol et al. [1] present an application of ERL-based classification to cherry fruit phenology prediction, comparing CORE, DMEL, and OCEC on a pomological dataset to derive interpretable rules for identifying the period most relevant to cherry fruit fly management. This illustrates how ERL can support explainable agricultural decision-making in a real-world classification task.

Alshaikh and Hewahi [2] present an educational application of LCSs in an intelligent tutoring system, where the task is to select an appropriate pedagogical tactic according to a student's performance. They compare standalone LCS and RL models with a hybrid model, PMCR, and report that the hybrid approach achieves the best accuracy on simulated student cases. Their approach shows how ERL can support adaptive tutoring decisions.

Sekine et al. [57] apply XCSI to en route arrival management, learning interpretable speed-control rules for inbound aircraft from a virtual design database generated by multiobjective optimization and an air-traffic simulator. In the Tokyo International Airport case study, the learned rules outperformed a decision-tree baseline and reduced both total flight time and fuel consumption, while remaining analyzable at the level of sector-specific and cross-sector traffic patterns.

Pour-Hosseini et al. [48] extended their earlier TinyXCS [49] by adding TinyDT, an offline decision-tree model trained from adaptive state-action data, and by framing the overall method as a two-model TinyML solution for resource-constrained cloud-edge environments such as IoT, autonomous systems, and industrial settings. In this framework, TinyXCS performs online rule-based workload allocation, deciding the portion of incoming tasks to process locally at the edge and the remainder to offload upward under latency, throughput, and power constraints.

Djartov et al. [11, 12, 13] apply an LCS to aviation decision support, particularly the dynamic alternate airport selection problem in time-critical in-flight emergencies. In [11], they introduced an interpretable LCS-based AI core for the Intelligent Pilot Advisory System (IPAS), framing alternate airport selection as a multi-criteria decision-making task. In [12], they extended this line by training an XCSF-style LCS on pilot survey data and analysing how pilots weigh factors such as landing distance, runway condition, and wind-related variables. These studies were later positioned within the broader IPAS framework [13].

Maurer et al. [35] apply an LCT-based lightweight rule-based RL controller to QoS-aware dynamic frequency scaling in mixed-critical embedded MPSoCs, where the agent is implemented on FPGA to react to sub-millisecond workload variations under power constraints. Their shielded SARSA(λ)-based design aims to maintain deadline-related QoS while reducing energy consumption through hardware-level runtime frequency control [35].

Islem Adnane and Zerari [26] apply ERL to classification rule mining in a business intelligence setting, with the aim of generating interpretable rules from high-dimensional and imbalanced data. They employ a quantum-inspired GA to optimize both attribute selection and categorical value representation under a Michigan-style rule representation, thereby improving the readability of the extracted classification rules while addressing dimensionality and discrete feature handling.

11 Research Tendencies and Identified Gaps

Over the survey period, research in ERL has been multi-faceted with progress in a multitude of areas as our well-populated seven categories indicate. We assess this multifaceted research landscape as healthy. Nonetheless, we observed the following tendencies.

One continued tendency of ERL research is moving beyond float-based or binary vectors as input such as time-to-event data [83, 84], text-data [73], or image data [51, 69, 87, 88]. Large population sizes remain a problem for ERL. Two-phase LCSs for drastically decreased population sizes show promise, as the novel HEROS [30] and continued research on SupRB [21] including a completed doctoral thesis [19] have shown. Yet, we feel that two-phase LCSs are not fully established in the field so far. Several ERL approaches have focused on resource-constrained TinyML or IoT applications [16, 35, 48, 49, 72], showing increased interest in this area in the community. Employing ERL for RL remains a topic of interest as several new works on ACS2 [18, 28, 41], XCS variants [51, 52, 53, 54, 78, 79] and other ERL systems [36, 58, 75] indicate.

Beyond this, we have identified further research gaps: ERL systems, in particular LCSs, are seen as interpretable due to their interpretable knowledge bases and decision-making. Yet, there remains room to increase the transparency via systemic approaches benefiting from the inherent interpretability.

Scalability remains a challenge for ERL systems, especially with large inputs and datasets. This is partly due to comparatively high computational cost. We see a need for more efficient systems or improved implementations. One option could be a higher degree of parallelization, as e.g. in AERL [59], see Section 5. Lower computational cost would allow more use-cases motivating more widespread adoption. Rule update schemes are one cost factor, as they are often expensive and stochastic. Developing more efficient update methods could improve both scalability and learning performance.

Many ERL approaches target single-objective problems, but extending them to multi-objective optimization could improve real-world effectiveness and efficiency. One system reacting to this need is e.g. Survival-LCS [83, 84], see Section 8.2.

Lastly, hybrid systems that combine ERL with other ML approaches such as deep learning, see Section 7, or other symbolic AI, promise to combine the benefits of different "worlds", creating more flexible, powerful systems.

12 Conclusion

This paper provides an overview of the published research in ERL between 11 March 2024 and 03 February 2026, including both dates. We categorise 63 examined papers into seven categories and provide a short summary of each of them. Only three contributions fall under Theory & Meta analysis, where especially theoretical research seems to have dwindled recently. There are seven papers to systems with new ERL architectures, indicating a high degree of novel ideas in the field. Also, Fuzzy-ERL received attention with five contributions. Research in deep learning and ERL hybrid systems led to five publications with substantially different methodologies. Furthermore, research in ERL has been particularly multi-faceted as 21 publications indicate contributing improvements and extensions to XCS, UCS and ExSTraCS 2.0, ACS2, SupRB, the Ant-Miner, and other systems. Further three publications focused on benchmarking ERL algorithms, while a total of 14 publications contributed application-based research demonstrating a high interest in use-cases for ERL systems. We identified a few further tendencies in the field and remaining research gaps. Overall, the research landscape of ERL appears to be multi-faceted and healthy.

References

[1] Erhan Akyol, Bilal Alatas, and Inanc Ozgen. 2025. Evolutionary algorithm approaches for cherry fruit classification based on pomological features. *Agriculture*, 15, 21. doi:10.3390/agriculture15212207.

[2] Fatema Alshaikh and Nabil Hewahi. 2025. An adaptive pedagogical model based on the classifier system and reinforcement learning. *The Computer Journal*, 68,

11, 1827–1841. eprint: https://academic.oup.com/comjnl/article-pdf/68/11/1827/63537697/bxaf077.pdf. doi:10.1093/comjnl/bxaf077.
[3] 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, 115–140. doi:10.1162/evco_a_00351.
[4] Umair Ayub and Bushra Almas. 2025. ASAQ—Ant-Miner: optimized rule-based classifier. *Soft Computing*, 29, 1, 355–364. doi:https://doi.org/10.1007/s00500-025-10474-6.
[5] Jaume Bacardit and Natalio Krasnogor. 2009. A mixed discrete-continuous attribute list representation for large scale classification domains. In *Proceedings of the 11th Annual Conference on Genetic and Evolutionary Computation* (GECCO '09). Association for Computing Machinery, New York, NY, USA, 1155–1162. ISBN: 978-1-60558-325-9. doi:10.1145/1569901.1570057.
[6] James Brookhouse, Ayah Helal, and Fernando Otero. 2025. An ensemble ant colony optimization algorithm with a hybrid pheromone model for learning rule lists. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO '25). Association for Computing Machinery, NH Malaga Hotel, Malaga, Spain, 1532–1539. ISBN: 9798400714658. doi:10.1145/3712256.3726427.
[7] Larry Bull and Toby O'Hara. 2002. Accuracy-based Neuro and Neuro-Fuzzy Classifier Systems, 7.
[8] Martin Butz. 2002. *Anticipatory Learning Classifier Systems*. Kluwer Academic Publishers, Boston/Dordrecht/London. ISBN: 978-1-4613-5290-7. doi:10.1007/978-1-4615-0891-5.
[9] Cameron Coombe, David Howard, and Will Browne. 2024. Learning classifier systems as a solver for the abstraction and reasoning corpus. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '24 Companion). Association for Computing Machinery, Melbourne, VIC, Australia, 1770–1778. ISBN: 979-8-4007-0495-6. doi:10.1145/3638530.3664157.
[10] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6, 2, 182–197. doi:10.1109/4235.996017.
[11] Boris Djartov, Sanaz Mostaghim, Anne Papenfuß, and Matthias Wies. 2024. A learning classifier system approach to time-critical decision-making in dynamic alternate airport selection. In *2024 IEEE Congress on Evolutionary Computation (CEC)*. 2024 IEEE Congress on Evolutionary Computation (CEC), 1–8. doi:10.1109/CEC60901.2024.10612016.
[12] Boris Djartov, Anne Papenfuß, and Matthias Wies. 2025. Wings of wisdom: learning from pilot decision data with interpretable ai models. In *HCI International 2024 – Late Breaking Papers*. Don Harris, Wen-Chin Li, and Heidi Krömker, (Eds.) Springer Nature Switzerland, Cham, 241–256. ISBN: 978-3-031-76824-8. doi:10.1007/978-3-031-76824-8_17.
[13] Boris Djartov and Jakob Würfel. 2025. Navigating Decisions in the Cockpit: The Intelligent Pilot Advisory System. In *The Design of Human-Centered Artificial Intelligence for the Workplace*. Constantinos K. Coursaris, Joerg Beringer, Pierre-Majorique Léger, and Burak Öz, (Eds.) Springer Nature Switzerland, Cham, 305–325. ISBN: 978-3-031-83512-4. doi:10.1007/978-3-031-83512-4_18.
[14] Jan Drugowitsch. 2007. *Learning Classifier Systems from First Principles : A Probabilistic Reformulation of Learning Classifier Systems from the Perspective of Machine Learning*. Ph.D. Dissertation. University of Bath, Somerset, UK. https://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.500684.
[15] Melanie Feist, Martin Breitbach, Heiko Trötsch, Christian Becker, and Christian Krupitzer. 2022. Rango: an intuitive rule language for learning classifier systems in cyber-physical systems. In *2022 IEEE International Conference on Autonomic Computing and Self-Organizing Systems (ACSOS)*, 31–40. doi:10.1109/ACSOS55765.2022.00021.
[16] Melanie Feist, Uwe Brinkschulte, and Mathias Pacher. 2024. The Adaptation Mechanism of Chameleon - A Comprehensive Adaptive Middleware for Mixed-Critical Cyber-Physical Networks. In *2024 IEEE 27th International Symposium on Real-Time Distributed Computing (ISORC)*. 2024 IEEE 27th International Symposium on Real-Time Distributed Computing (ISORC), 1–10. doi:10.1109/ISORC61049.2024.10551345.
[17] Ghofrane Ben Hamouda, Lilia Rejeb, and Lamjed Ben Said. 2024. Ensemble learning for multi-channel sleep stage classification. *Biomedical Signal Processing and Control*, 93, 106184. doi:https://doi.org/10.1016/j.bspc.2024.106184.
[18] Tomohiro Hayashida, Shinya Sekizaki, Yushi Furuya, and Ichiro Nishizaki. 2025. ACS2-powered pedestrian flow simulation for crowd dynamics. *AppliedMath*, 5, 3. doi:10.3390/appliedmath5030088.
[19] Michael Heider. 2025. *Disentangling the Model Selection Tasks for Improved Explainability in a Rule-Based Machine Learning System*. doctoralthesis. Universität Augsburg, 206. Retrieved Feb. 3, 2026 from https://opus.bibliothek.uni-augsburg.de/opus4/frontdoor/index/index/docId/120504.
[20] Michael Heider. 2025. How to design an lcs to create explainable ai models for real-world applications. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '25 Companion). Association for Computing Machinery, NH Malaga Hotel, Malaga, Spain, 2258–2260. ISBN: 979-8-4007-1464-1. doi:10.1145/3712255.3734305.
[21] Michael Heider, Maximilian Krischan, Roman Sraj, and Jörg Hähner. 2024. Exploring self-adaptive genetic algorithms to combine compact sets of rules. In *2024 IEEE Congress on Evolutionary Computation (CEC)*. 2024 IEEE Congress on Evolutionary Computation (CEC). doi:10.1109/CEC60901.2024.10612101.
[22] Michael Heider, David Pätzel, Helena Stegherr, and Jörg Hähner. 2023. A Metaheuristic Perspective on Learning Classifier Systems. In *Metaheuristics for Machine Learning*. Mansour Eddaly, Bassem Jarboui, and Patrick Siarry, (Eds.) Springer Nature Singapore, Singapore, 73–98. ISBN: 978-981-19-3888-7. doi:10.1007/978-981-19-3888-7_3.
[23] Michael Heider, David Pätzel, and Alexander R. M. Wagner. 2022. An overview of LCS research from 2021 to 2022. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 2086–2094. doi:10.1145/3520304.3533985.
[24] Michael Heider, Helena Stegherr, Richard Nordsieck, and Jörg Hähner. 2023. Assessing Model Requirements for Explainable AI: A Template and Exemplary Case Study. *Artificial Life*, 29, 4, 468–486. doi:10.1162/artl_a_00414.
[25] Michael Heider, Helena Stegherr, Roman Sraj, David Pätzel, Jonathan Wurth, and Jörg Hähner. 2023. SupRB in the context of rule-based machine learning methods: a comparative study. *Applied Soft Computing*, 147, 110706. doi:https://doi.org/10.1016/j.asoc.2023.110706.
[26] Youcef Islem Adnane and Mounira Zerari. 2024. Optimizing business intelligence classification rule mining using quantum-inspired genetic algorithm. *IEEE Access*, 12, 137041–137053. doi:10.1109/ACCESS.2024.3463506.
[27] Konstantinos Kovas and Ioannis Hatzilygeroudis. 2024. ACRES: a framework for (semi)automatic generation of rule-based expert systems with uncertainty from datasets. *Expert Systems*, 41, 12. doi:https://doi.org/10.1111/exsy.13723.
[28] Mateusz Łabędzki and Olgierd Unold. 2025. Why state differentiation in ACS2 is not enough in aliased environments. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '25 Companion). Association for Computing Machinery, New York, NY, USA, 2249–2257. ISBN: 979-8-4007-1464-1. doi:10.1145/3712255.3734296.
[29] Pier Luca Lanzi and Daniele Loiacono. 2025. A proposal for a leaner narrative of learning classifier systems. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '25 Companion). Association for Computing Machinery, NH Malaga Hotel, Malaga, Spain, 2261–2268. ISBN: 979-8-4007-1464-1. doi:10.1145/3712255.3735661.
[30] Gabriel Lipschutz-Villa, Harsh Bandhey, Ruonan Yin, Malek Kamoun, and Ryan Urbanowicz. 2025. Rule-based machine learning: separating rule and rule-set pareto-optimization for interpretable noise-agnostic modeling. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO '25). Association for Computing Machinery, NH Malaga Hotel, Malaga, Spain, 407–415. ISBN: 979-8-4007-1465-8. doi:10.1145/3712256.3726461.
[31] Yi Liu, Ying Bi, Yu Cui, Chengyuan Zhu, Wen Cheng, Mingkai Sheng, and Yiding Zhang. 2024. Adapting explainable Decision Sets to Continuous-based Real-World Problems. In *2024 7th International Symposium on Autonomous Systems (ISAS)*. 2024 7th International Symposium on Autonomous Systems (ISAS), 1–6. doi:10.1109/ISAS61044.2024.10552551.
[32] Yi Liu, Yu Cui, Will Browne, Bing Xue, Wen Cheng, Yong Li, and Lingfang Zeng. 2023. Absumption and Subsumption based Learning Classifier System for Real-World Continuous-based Problems. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*. GECCO '23 Companion: Companion Conference on Genetic and Evolutionary Computation. ACM, Lisbon Portugal, 299–302. ISBN: 979-8-4007-0120-7. doi:10.1145/3583133.3590564.
[33] Yi Liu, Yu Cui, Wen Cheng, Will Neil Browne, Bing Xue, Chengyuan Zhu, Yiding Zhang, Mingkai Sheng, and Lingfang Zeng. 2024. A phenotypic learning classifier system for problems with continuous features. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO '24). Association for Computing Machinery, Melbourne, VIC, Australia, 349–357. ISBN: 979-8-4007-0494-9. doi:10.1145/3638529.3654007.
[34] Zafar Masood, Zheng Jiangbin, Idrees Ahmad, Chai Dongdong, Wasif Shabbir, and Muhammad Irfan. 2024. A novel continual reinforcement learning-based expert system for self-optimization of soft real-time systems. *Expert Systems with Applications*, 238, 122309. doi:https://doi.org/10.1016/j.eswa.2023.122309.
[35] Florian Maurer, Michael Meidinger, Yiming Lu, Matthias Schlemmer, Thomas Hallermeier, Anmol Surhonne, Thomas Wild, and Andreas Herkersdorf. 2025. Rule-Based Reinforcement Learning on FPGA for QoS-Aware Dynamic Frequency Scaling. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 33, 11, 3008–3017. doi:10.1109/TVLSI.2025.3592300.
[36] Matthew Mitchell. 2024. A hybrid connectionist/LCS for hidden-state problems. *Neural Computing and Applications*, 36, 22, 13579–13603. doi:10.1007/s00521-024-09758-z.
[37] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing Atari with Deep Reinforcement Learning. *CoRR*, abs/1312.5602, 9. http://arxiv.org/abs/1312.5602.
[38] Mohd Mizan Munif, Husna binti Jamal Abdul Nasir, and Muhammad Imran Ahmad. 2024. Enhanced ant colony system for water quality rules classification in smart aquaculture. In *2024 IEEE Symposium on Wireless Technology & Applications (ISWTA)*, 216–220. doi:10.1109/ISWTA62130.2024.10652089.
[39] Muhammad Fawad Nasim, Muhammad Anwar, Almuhannad S Alorfi, Hamza Awad Ibrahim, Ali Ahmed, Arfan Jaffar, Sheeraz Akram, Abubakar Siddique,

and Hafiz Muhammad Zeeshan. 2025. Cognitively inspired sound-based automobile problem detection: a step toward explainable AI (XAI). *International Journal of Advanced and Applied Sciences*, 12, 8, 1–15. doi:10.21833/ijaas.2025.08.001.

[40] Damijan Novak, Domen Verber, Iztok Fister, and Iztok Fister. 2024. Association rule mining as knowledge infusion into the mechanics of an extended classifier system. In *2024 IEEE 28th International Conference on Intelligent Engineering Systems (INES)*. 2024 IEEE 28th International Conference on Intelligent Engineering Systems (INES), 000203–000208. doi:10.1109/INES63318.2024.10629091.

[41] Romain Orhand and Ali Ayadi. 2025. Towards a neurosymbolic approach based on anticipatory learning classifier systems and ontologies. *Procedia Computer Science*, 270, 1232–1241. 29th International Conference on Knowledge-Based and Intelligent Information & Engineering Systems (KES 2025). doi:https://doi.org/10.1016/j.procs.2025.09.244.

[42] Rafael S. Parpinelli, Heitor S. Lopes, and Alex A. Freitas. 2002. Data mining with an ant colony optimization algorithm. *IEEE Transactions on Evolutionary Computation*, 6, 4, 321–332. doi:10.1109/TEVC.2002.802452.

[43] David Pätzel, Michael Heider, and Alexander R. M. Wagner. 2021. An overview of LCS research from 2020 to 2021. In *GECCO '21: Genetic and Evolutionary Computation Conference, Companion Volume, Lille, France, July 10-14, 2021*. Krzysztof Krawiec, (Ed.) ACM, 1648–1656. doi:10.1145/3449726.3463173.

[44] David Pätzel, Richard Nordsieck, and Jörg Hähner. 2024. A closer look at length-niching selection and spatial crossover in variable-length evolutionary rule set learning. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '24 Companion). Association for Computing Machinery, Melbourne, VIC, Australia, 1779–1787. ISBN: 979-8-4007-0495-6. doi:10.1145/3638530.3664178.

[45] David Pätzel, Richard Nordsieck, and Jörg Hähner. 2024. Length-niching selection and spatial crossover in variable-length evolutionary rule set learning. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '24 Companion). Association for Computing Machinery, Melbourne, VIC, Australia, 303–306. ISBN: 9798400704956. doi:10.1145/3638530.3654308.

[46] David Pätzel, Richard Nordsieck, and Jörg Hähner. 2024. Measuring similarities in model structure of metaheuristic rule set learners. In *Applications of Evolutionary Computation*. Stephen Smith, João Correia, and Christian Cintrano, (Eds.) Springer Nature Switzerland, Cham, 256–272. ISBN: 978-3-031-56855-8. doi:10.1007/978-3-031-56855-8_16.

[47] David Pätzel, Anthony Stein, and Masaya Nakata. 2020. An overview of LCS research from IWLCS 2019 to 2020. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion* (GECCO '20). Association for Computing Machinery, Cancún, Mexico, 1782–1788. ISBN: 9781450371278. doi:10.1145/3377929.3398105.

[48] Mohammad R Pour-Hosseini, Mahdi Abbasi, Atefeh Salimi, Erik Elmroth, Hassan Haghighi, Parham Moradi, and Bahman Javadi. 2025. Tiny machine learning models for autonomous workload distribution across cloud-edge computing continuum. *Cluster Computing*, 28, 6, 381. doi:10.1007/s10586-025-05289-x.

[49] Mohammad Reza Pour-Hosseini, Mahdi Abbasi, Bahman Javadi, Seifeddine Benelghali, Hassan Haghighi, Ehsan Mohammadi-Pasand, and Parham Moradi. 2024. Efficient Edge Computing: Harnessing Compact Machine Learning Models for Workload Optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '24 Companion). Association for Computing Machinery, New York, NY, USA, 239–242. ISBN: 979-8-4007-0495-6. doi:10.1145/3638530.3654307.

[50] Richard John Preen, Stewart W. Wilson, and Larry Bull. 2021. Autoencoding with a Classifier System. *IEEE Transactions on Evolutionary Computation*, PP, 1–1. doi:10.1109/TEVC.2021.3079320.

[51] Connor Schönberner, Armin Mackensen, and Sven Tomforde. 2025. Dimensionality reduction for enabling visual reinforcement learning with a classifier system. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '25 Companion). Association for Computing Machinery, NH Malaga Hotel, Malaga, Spain, 2269–2278. ISBN: 9798400714641. doi:10.1145/3712255.3734321.

[52] Connor Schönberner, Kjell-Matti Rothenburger, Armin Mackensen, and Sven Tomforde. 2026. Neural rules for reinforcement learning with XCSF as controller in organic computing systems. In *Architecture of Computing Systems*. Sven Tomforde, Christian Krupitzer, Stéphane Vialle, Estela Suarez, and Thilo Pionteck, (Eds.) Springer Nature Switzerland, Cham, 268–282. ISBN: 978-3-032-03281-2. doi:10.1007/978-3-032-03281-2_18.

[53] Connor Schönberner and Sven Tomforde. 2025. Double XCSF on target? *Evolutionary Computation*, 1–36. eprint: https://direct.mit.edu/evco/article-pdf/doi/10.1162/EVCO.a.377/2571560/evco.a.377.pdf. doi:10.1162/EVCO.a.377.

[54] Connor Schönberner and Sven Tomforde. 2024. XCS: Is Covering All You Need? In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '24 Companion). Association for Computing Machinery, New York, NY, USA, 1788–1796. ISBN: 979-8-4007-0495-6. doi:10.1145/3638530.3664146.

[55] Gijs Schröder and Johannes Textor. 2025. Classifier systems as linear probability models. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO '25). Association for Computing Machinery, NH Malaga Hotel, Malaga, Spain, 167–175. ISBN: 9798400714658. doi:10.1145/3712256.3726453.

[56] Gijs Schröder, Inge Wortel, and Johannes Textor. 2024. Population-based algorithms built on weighted automata. In *Parallel Problem Solving from Nature – PPSN XVIII*. Michael Affenzeller, Stephan M. Winkler, Anna V. Kononova, Heike Trautmann, Tea Tušar, Penousal Machado, and Thomas Bäck, (Eds.) Springer Nature Switzerland, Cham, 315–332. ISBN: 978-3-031-70071-2.

[57] Katsuhiro Sekine, Tomoaki Tatsukawa, Takeru Aoki, Kozo Fujii, and Eri Itoh. 2025. Adaptive and interpretable speed control rules for en route arrival management using extended learning classifier system. *IEEE Access*, 13, 39631–39649. doi:10.1109/ACCESS.2025.3546640.

[58] Hormoz Shahrzad, Babak Hodjat, and Risto Miikkulainen. 2025. EVOTER: evolution of transparent explainable rule-sets. *ACM Trans. Evol. Learn. Optim.*, 5, 2, Article 11, 30 pages. doi:10.1145/3702651.

[59] Hormoz Shahrzad and Risto Miikkulainen. 2025. GPU-accelerated rule evaluation and evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '25 Companion). Association for Computing Machinery, NH Malaga Hotel, Malaga, Spain, 2279–2281. ISBN: 979-8-4007-1464-1. doi:10.1145/3712255.3736397.

[60] Hiroki Shiraishi, Yohei Hayamizu, and Tomonori Hashiyama. 2023. Fuzzy-UCS revisited: self-adaptation of rule representations in michigan-style learning fuzzy-classifier systems. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO '23). Association for Computing Machinery, Lisbon, Portugal, 548–557. ISBN: 9798400701191. doi:10.1145/3583131.3590360.

[61] Hiroki Shiraishi, Yohei Hayamizu, Tomonori Hashiyama, Keiki Takadama, Hisao Ishibuchi, and Masaya Nakata. 2026. Adapting rule representation with four-parameter beta distribution for learning classifier systems. *IEEE Transactions on Evolutionary Computation*, 30, 1, 363–377. doi:10.1109/TEVC.2025.3550915.

[62] Hiroki Shiraishi, Yohei Hayamizu, Tomonori Hashiyama, Keiki Takadama, Hisao Ishibuchi, and Masaya Nakata. 2025. Evolutionary co-optimization of rule shape and fuzziness in rule-based machine learning. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '25 Companion). Association for Computing Machinery, NH Malaga Hotel, Malaga, Spain, 71–72. ISBN: 9798400714641. doi:10.1145/3712255.3734248.

[63] Hiroki Shiraishi, Hisao Ishibuchi, and Masaya Nakata. 2026. A class inference scheme with dempster–shafer theory for learning fuzzy-classifier systems. *ACM Trans. Evol. Learn. Optim.*, 6, 1, Article 4, 4:1–4:43. doi:10.1145/3717613.

[64] Hiroki Shiraishi, Hisao Ishibuchi, and Masaya Nakata. 2025. Evidential fuzzy rule-based machine learning to quantify classification uncertainty. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '25 Companion). Association for Computing Machinery, NH Malaga Hotel, Malaga, Spain, 69–70. ISBN: 9798400714641. doi:10.1145/3712255.3734234.

[65] Hiroki Shiraishi, Hisao Ishibuchi, and Masaya Nakata. 2025. Stochastic local mutation for efficient membership function exploration in fuzzy classifiers. In *2025 IEEE International Conference on Fuzzy Systems (FUZZ)*. 2025 IEEE International Conference on Fuzzy Systems (FUZZ), 1–6. doi:10.1109/FUZZ62266.2025.11152164.

[66] Hiroki Shiraishi, Hisao Ishibuchi, and Masaya Nakata. 2025. X-KAN: optimizing local kolmogorov-arnold networks via evolutionary rule-based machine learning. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*. James Kwok, (Ed.) Main Track. International Joint Conferences on Artificial Intelligence Organization, 8930–8938. arXiv: 2505.14273 [cs]. doi:10.24963/ijcai.2025/993.

[67] Hiroki Shiraishi, Rongguang Ye, Hisao Ishibuchi, and Masaya Nakata. 2024. A variable-length fuzzy set representation for learning fuzzy-classifier systems. In *Parallel Problem Solving from Nature – PPSN XVIII*. Michael Affenzeller, Stephan M. Winkler, Anna V. Kononova, Heike Trautmann, Tea Tušar, Penousal Machado, and Thomas Bäck, (Eds.) Springer Nature Switzerland, Cham, 386–402. ISBN: 978-3-031-70071-2. doi:10.1007/978-3-031-70071-2_24.

[68] 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* (GECCO '24 Companion). Association for Computing Machinery, Melbourne, VIC, Australia, 1797–1806. ISBN: 9798400704956. doi:10.1145/3638530.3664165.

[69] Abubakar Siddique, Muhammad Iqbal, and Will N. Browne. 2024. Enhancing visual focus: incorporating attention into rule-based machine learning. In *2024 39th International Conference on Image and Vision Computing New Zealand (IVCNZ)*. 2024 39th International Conference on Image and Vision Computing New Zealand (IVCNZ), 1–6. doi:10.1109/IVCNZ64857.2024.10794461.

[70] Anthony Stein, Roland Maier, Lukas Rosenbauer, and Jörg Hähner. 2020. XCS classifier system with experience replay. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*. GECCO '20: Genetic and Evolutionary Computation Conference. ACM, Cancún Mexico, 404–413. ISBN: 978-1-4503-7128-5. doi:10.1145/3377930.3390249.

[71] Marco Steinberger, Roman Küble, Michael Heider, and Jörg Hähner. 2026. Evaluating adaptive systems: a comparative study of xcs and established reinforcement learning algorithms in noisy multi-step environments. In *38th*

GI/ITG International Conference on Architecture of Computing Systems (ARCS). Sven Tomforde, Christian Krupitzer, Stéphane Vialle, Estela Suarez, and Thilo Pionteck, (Eds.) Springer Nature Switzerland, Cham, 283–297. ISBN: 978-3-032-03281-2.

[72] Anmol Surhonne, Manuel Wensauer, Florian Maurer, Thomas Wild, and Andreas Herkersdorf. 2024. XCS with dynamic sized experience replay for memory constrained applications. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '24 Companion). Association for Computing Machinery, Melbourne, VIC, Australia, 1807–1814. ISBN: 979-8-4007-0495-6. doi:10.1145/3638530.3664148.

[73] Yossathorn Tianrungroj and Hitoshi Iba. 2024. ROIL: rule optimization via large language model for imitation learning. In *2024 IEEE Congress on Evolutionary Computation (CEC)*. doi:10.1109/CEC60901.2024.10612166.

[74] Michael C. Todd and Gilbert L. Peterson. 2024. Temporal metadata analysis: A learning classifier system approach. *Forensic Science International: Digital Investigation*, 51, 301842. doi:https://doi.org/10.1016/j.fsidi.2024.301842.

[75] Sz-Ting Tzeng, Nirav Ajmeri, and Munindar P. Singh. 2024. Value-based rationales improve social experience: A multiagent simulation study. In *ECAI 2024 - 27th European Conference on Artificial Intelligence, 19-24 October 2024, Santiago de Compostela, Spain - Including 13th Conference on Prestigious Applications of Intelligent Systems (PAIS 2024)* (Frontiers in Artificial Intelligence and Applications). Ulle Endriss, Francisco S. Melo, Kerstin Bach, Alberto José Bugarín Diz, Jose Maria Alonso-Moral, Senén Barro, and Fredrik Heintz, (Eds.) Vol. 392. IOS Press, 3612–3619. doi:10.3233/FAIA240917.

[76] Ryan Urbanowicz, Gediminas Bertasius, and Jason Moore. 2014. An Extended Michigan-Style Learning Classifier System for Flexible Supervised Learning, Classification, and Data Mining. In *Parallel Problem Solving from Nature - PPSN XIII* number 8672. 13th International Conference on Parallel Problem Solving from Nature. Vol. Lecture Notes in Computer Science, vol 8672. Springer, Cham. Springer, Cham, 211–221. ISBN: 978-3-319-10761-5. doi:10.1007/978-3-319-10762-2_21.

[77] Ryan J. Urbanowicz and Jason H. Moore. 2009. Learning Classifier Systems: A Complete Introduction, Review, and Roadmap. *Journal of Artificial Evolution and Applications*.

[78] Fumito Uwano and Will N. Browne. 2025. Enhancing XCS with dual-stream identification for perceptual aliasing in multi-step decision-making. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO '25). Association for Computing Machinery, NH Malaga Hotel, Malaga, Spain, 499–507. ISBN: 9798400714658. doi:10.1145/3712256.3726470.

[79] Fumito Uwano and Will Neil Browne. 2024. Cognitive learning system for sequential aliasing patterns of states in multistep decision-making. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '24 Companion). Association for Computing Machinery, Melbourne, VIC, Australia, 315–318. ISBN: 979-8-4007-0495-6. doi:10.1145/3638530.3654110.

[80] Hado Van Hasselt, Arthur Guez, and David Silver. 2016. Deep Reinforcement Learning with Double Q-Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30, 1. doi:10.1609/aaai.v30i1.10295.

[81] Hado van Hasselt. 2010. Double Q-learning. In *Advances in Neural Information Processing Systems*. J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta, (Eds.) Vol. 23. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2010/file/091d584fced301b442654dd8c23b3fc9-Paper.pdf.

[82] Stewart W. Wilson. 1995. Classifier Fitness Based on Accuracy. *Evolutionary Computation*, 3, 2, 149–175. doi:10.1162/evco.1995.3.2.149.

[83] Alexa Woodward, Harsh Bandhey, Jason H. Moore, and Ryan J. Urbanowicz. 2024. Survival-LCS: a rule-based machine learning approach to survival analysis. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO '24). Association for Computing Machinery, Melbourne, VIC, Australia, 431–439. ISBN: 9798400704949. doi:10.1145/3638529.3654154.

[84] Alexa A. Woodward, Harsh Bandhey, Jason H. Moore, and Ryan J. Urbanowicz. 2025. Survival-LCS: rule-based survival analysis without proportional hazard assumptions. *ACM Trans. Evol. Learn. Optim.* Just Accepted. doi:10.1145/3772719.

[85] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. ReAct: synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*. https://openreview.net/forum?id=WE_vluYUL-X.

[86] Naoya Yatsu, Hiroki Shiraishi, Hiroyuki Sato, and Keiki Takadama. 2023. Exploring high-dimensional rules indirectly via latent space through a dimensionality reduction for XCS. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO '23). Association for Computing Machinery, Lisbon, Portugal, 606–614. ISBN: 9798400701191. doi:10.1145/3583131.3590439.

[87] Naoya Yatsu, Hiroki Shiraishi, Hiroyuki Sato, and Keiki Takadama. 2024. Generating high-dimensional prototypes with a classifier system by evolving in latent space. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion* (GECCO '24 Companion). Association for Computing Machinery, Melbourne, VIC, Australia, 2127–2130. ISBN: 9798400704956. doi:10.1145/3638530.3654248.

[88] Naoya Yatsu, Hiroki Shiraishi, Hiroyuki Sato, and Keiki Takadama. 2024. Prototype generation with the supervised classifier system on kNN matching. In *2024 IEEE Congress on Evolutionary Computation (CEC)*. 2024 IEEE Congress on Evolutionary Computation (CEC), 1–8. doi:10.1109/CEC60901.2024.10611796.

[89] Zhaoxiang Zang, Zhao Li, Zhiping Dan, and Junying Wang. 2024. Improving selection strategies in zeroth-level classifier systems based on average reward reinforcement learning. *Journal of Ambient Intelligence and Humanized Computing*, 15, 2, 1201–1211. doi:10.1007/s12652-018-0682-x.

[90] Ryan Zhou et al. 2024. Evolutionary computation and explainable AI: a roadmap to understandable intelligent systems. *IEEE Transactions on Evolutionary Computation*, 29, 5, 2213–2228. doi:10.1109/TEVC.2024.3476443.