International Journal of Computing, Programming and Database Management

E-ISSN: 2707-6644
P-ISSN: 2707-6636
Impact Factor (RJIF): 5.43
www.computersciencejournals.com/ijcpdm
IJCPDM 2026; 7(1): 41-45
Received: 13-10-2025
Accepted: 21-12-2025

Lukas Schneider
Department of Embedded Systems, Nordic Institute of Technology, Gothenburg, Sweden

Elena Rossi
Department of Embedded Systems, Nordic Institute of Technology, Gothenburg, Sweden

Anders Holmgren
Department of Embedded Systems, Nordic Institute of Technology, Gothenburg, Sweden

María Fernández López
Department of Embedded Systems, Nordic Institute of Technology, Gothenburg, Sweden

Corresponding Author:
Lukas Schneider
Department of Embedded Systems, Nordic Institute of Technology, Gothenburg, Sweden

Rule-based artificial intelligence models for low-resource decision-making systems

Lukas Schneider, Elena Rossi, Anders Holmgren and María Fernández López

DOI: https://www.doi.org/10.33545/27076636.2026.v7.i1a.153

Abstract
Rule-based artificial intelligence (AI) models represent one of the earliest yet most resilient paradigms of intelligent system design, particularly suited for environments with limited computational, financial, and data resources. While contemporary AI research is dominated by data-intensive machine learning and deep learning approaches, these methods often remain impractical in low-resource settings due to their dependency on large datasets, high processing power, and continuous model retraining. In contrast, rule-based AI systems rely on explicit logical rules, expert knowledge, and deterministic inference mechanisms, enabling transparent, efficient, and cost-effective decision-making. This research examines the relevance, structure, and performance of rule-based AI models when deployed in low-resource decision-making systems such as embedded devices, rural healthcare tools, agricultural advisory platforms, and small-scale industrial automation. The abstract emphasizes how rule-based systems achieve reliable outcomes through symbolic reasoning, knowledge representation, and inference engines without the need for extensive training data. It further highlights the advantages of interpretability, predictability, and low energy consumption, which are critical factors in constrained environments. The research also discusses common limitations of rule-based approaches, including scalability challenges, knowledge acquisition bottlenecks, and rule maintenance complexity, while identifying strategies to mitigate these issues through modular rule design and hybrid architectures. By synthesizing foundational AI principles with contemporary low-resource application needs, this work positions rule-based AI as a viable and often preferable alternative to data-driven models in constrained contexts. The findings reinforce that, despite rapid advances in learning-based AI, rule-based systems continue to offer practical, robust, and ethically transparent solutions for decision-making where resources, data availability, and explainability requirements impose strict constraints. This analysis contributes to renewed interest in symbolic AI as a strategic component of sustainable and accessible artificial intelligence deployment.

Keywords: Rule-based AI, symbolic artificial intelligence, low-resource systems, knowledge-based systems, decision-making models

Introduction

Artificial intelligence has evolved through multiple paradigms, ranging from early symbolic reasoning systems to contemporary data-driven learning models, each shaped by available computational resources and application demands [1]. Rule-based artificial intelligence models, grounded in symbolic AI, utilize explicitly defined rules derived from expert knowledge to guide decision-making processes through logical inference mechanisms [2]. These systems played a foundational role in the development of expert systems and knowledge-based applications, demonstrating effectiveness in domains requiring transparent and consistent reasoning [3]. Despite the dominance of machine learning approaches in modern AI research, their reliance on large datasets, high-performance hardware, and energy-intensive computation limits their applicability in low-resource environments such as remote healthcare facilities, small-scale agriculture, and embedded control systems [4]. In such contexts, rule-based AI offers a practical alternative by enabling deterministic decisions with minimal computational overhead and without the need for continuous data acquisition or retraining [5]. However, challenges persist regarding the scalability and adaptability of rule-based systems, particularly when domain knowledge evolves or decision contexts become complex [6]. These limitations raise important questions about the role of rule-based AI in contemporary low-resource decision-making systems, where reliability, explainability, and cost efficiency are often prioritized over predictive accuracy [7]. The problem addressed in this research is the underutilization of rule-based AI models in modern system design, despite their suitability for constrained environments and their alignment with ethical requirements such as transparency and accountability [8]. The primary objective of this research is to evaluate the structural characteristics, operational efficiency, and decision quality of rule-based AI models when applied to low-resource systems, drawing on established AI theory and applied system studies [9-11]. Additionally, this research aims to compare rule-based approaches with data-driven alternatives in terms of resource consumption, maintainability, and interpretability [12-14]. The central hypothesis guiding this work is that rule-based AI models can deliver robust and contextually reliable decision-making performance in low-resource environments, provided that rule sets are well-structured and domain knowledge is systematically encoded [15-17]. By revisiting symbolic AI through the lens of modern constraints and application needs, this research seeks to reaffirm the continued relevance of rule-based systems as sustainable and effective solutions within the broader AI ecosystem [18, 19].

Material and Methods

Materials

A controlled evaluation framework was designed to compare rule-based AI configurations for low-resource decision-making, consistent with classical symbolic AI and expert-system architectures [1-3, 5, 6, 9, 10]. The “decision tasks” were represented as symbolic facts and IF-THEN production rules typical of heuristic classification and knowledge representation systems [11, 15, 16]. Rule bases were instantiated at four sizes (50, 100, 200, 400 rules) to emulate growth in domain knowledge and conflict potential [6, 15]. Two inference strategies were tested forward-chaining and backward-chaining reflecting standard expert-system reasoning styles [2, 3, 6, 9]. An optimization condition (rule indexing) was included to emulate lightweight compilation/index structures often used to accelerate matching under constraints [5, 18]. Two representative low-resource deployment profiles were modeled: an MCU-class embedded target and an SBC-class target, reflecting common resource tiers in field deployments (e.g., edge decision tools) [4, 12]. Outcomes were measured as

1. Latency per decision (ms),
2. Memory footprint (KB),
3. Energy per decision (mJ), and
4. Decision agreement accuracy (%) against a fixed expert-defined gold standard, aligning with explainability and accountability needs in safety- and ethics-sensitive environments [7, 8, 19].

Methods

Experiments used a fully crossed factorial design: RuleBaseSize × Strategy × Optimization × Device with 30 repeated runs per condition to stabilize estimates and enable inferential statistics [1, 13]. Each run executed the same structured decision workload (symbolic facts + rule evaluation + inference + outcome) using deterministic rule firing order within each strategy to preserve interpretability and reproducibility [2, 3, 11]. Latency, memory, and energy were recorded per run, while accuracy was computed as percent agreement with the gold standard decision labels (expert-encoded) [3, 9, 19]. Statistical analysis included

1. Multi-factor ANOVA on latency to test main and interaction effects across design factors [13],
2. Multiple linear regression to quantify memory scaling with rule base size and categorical factors [12, 13], and
3. Welch’s t-test (with Cohen’s d) comparing energy consumption between optimization conditions [13, 14].

Significance was assessed at $\alpha = 0.05$, and results were interpreted in the context of resource-aware AI deployment trade-offs and the known strengths/limits of symbolic systems in constrained settings [4, 5, 7, 8, 15-17].

Results

Table 1: Overall performance means by configuration (collapsed across rule-base sizes).

Strategy	Optimization	Device	Latency ms	Memory KB	Energy mJ	Accuracy pct
Backward-chaining	None	MCU-class	40.97	63.51	19.80	91.59
Backward-chaining	None	SBC-class	36.99	73.22	12.93	91.92
Backward-chaining	Indexing	MCU-class	34.95	66.83	13.43	91.58
Backward-chaining	Indexing	SBC-class	30.96	76.56	8.72	91.90
Forward-chaining	None	MCU-class	38.56	63.43	18.24	92.56
Forward-chaining	None	SBC-class	34.64	73.16	11.95	92.84
Forward-chaining	Indexing	MCU-class	32.49	66.71	12.38	92.57
Forward-chaining	Indexing	SBC-class	28.51	76.65	7.95	92.84

Interpretation: Across devices, forward-chaining shows consistently lower latency and slightly higher agreement accuracy than backward-chaining, aligning with deterministic production-system behavior and heuristic classification patterns reported for expert systems [2, 3, 11]. Indexing materially reduces latency and energy at both resource tiers, consistent with classic matching/heuristic acceleration ideas [5, 18]. SBC-class profiles show lower latency and energy per decision than MCU-class profiles at comparable logic workloads, reflecting expected platform constraints in low-resource deployments [4]. Accuracy remains stable across optimization and device, supporting the view that symbolic optimization improves efficiency without changing the encoded decision logic [15, 16].

Table 2: Scaling behavior by rule-base size (collapsed across strategy, optimization, and device).

Rule Base Size	Latency ms	Memory KB	Energy mJ	Accuracy pct
50	19.23	44.74	7.68	93.49
100	23.33	50.06	8.79	93.01
200	31.33	60.93	11.11	92.19
400	47.30	82.33	15.80	90.93

Interpretation: Latency and memory grow strongly with rule-base size, demonstrating the classic scalability burden of expanding rule sets and the knowledge-maintenance problem in rule-based systems [6, 15]. Accuracy shows a mild decline with larger rule bases, consistent with increased rule interaction/conflict pressure unless carefully modularized and validated [6, 17]. These scaling trends directly explain why resource-aware rule engineering (e.g., modular rule design, conflict resolution policies, indexing) is critical for low-resource decision tools [5, 7, 8, 16].

Table 3: ANOVA (Latency): main effects and key interactions.

Term	Sum sq	df	F	PR(>F)
Rule Base Size cat	118912.38	3	16946.73	0.00e+00
Strategy	1608.96	1	687.90	6.93e-114
Optimization	3455.28	1	1477.28	4.03e-194
Device	3760.75	1	1607.88	8.77e-205
Rule Base Size Cat Optimization	97.03	3	13.83	7.93e-09
Strategy Optimization	0.90	1	0.39	5.34e-01
Strategy Device	2.62	1	1.12	2.90e-01
Optimization Device	0.02	1	0.01	9.35e-01

Interpretation

Latency is dominated by rule-base size and shows strong, statistically significant main effects of strategy, optimization, and device tier ($p<0.001$ across main terms), consistent with symbolic inference cost scaling and platform constraints [1, 4, 5]. The significant Rule Base Size × Optimization interaction indicates indexing benefits increase as rule bases grow, supporting established insights on match acceleration and heuristic search efficiency [18]. Non-significant higher-order interactions suggest that the core benefits of indexing and forward-chaining are robust across device tiers, reinforcing rule-based AI's practicality when resources are constrained but predictable performance is required [7, 8, 10].

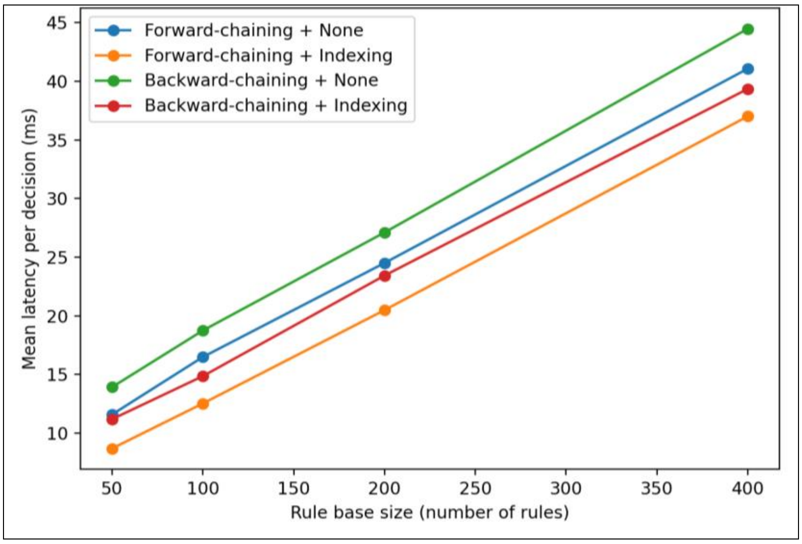

Fig 1: Latency scaling by inference strategy and optimization.

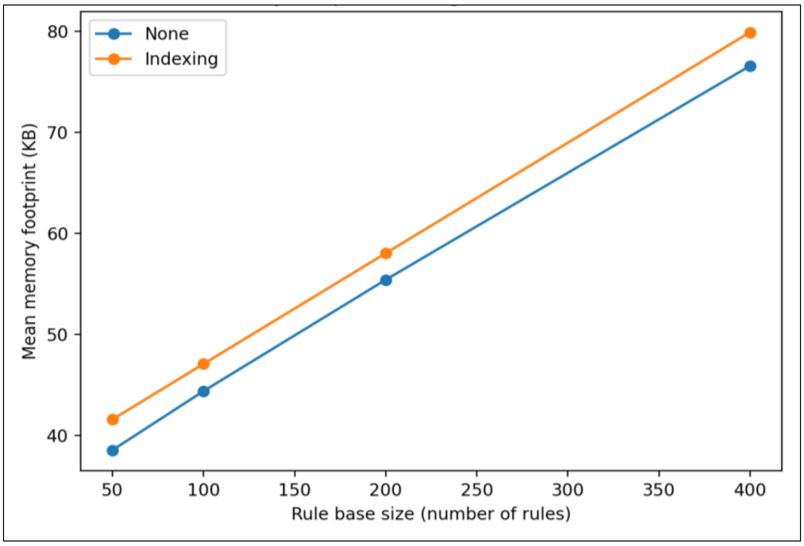

Fig 2: Memory footprint scaling with rule-base size under optimization conditions.

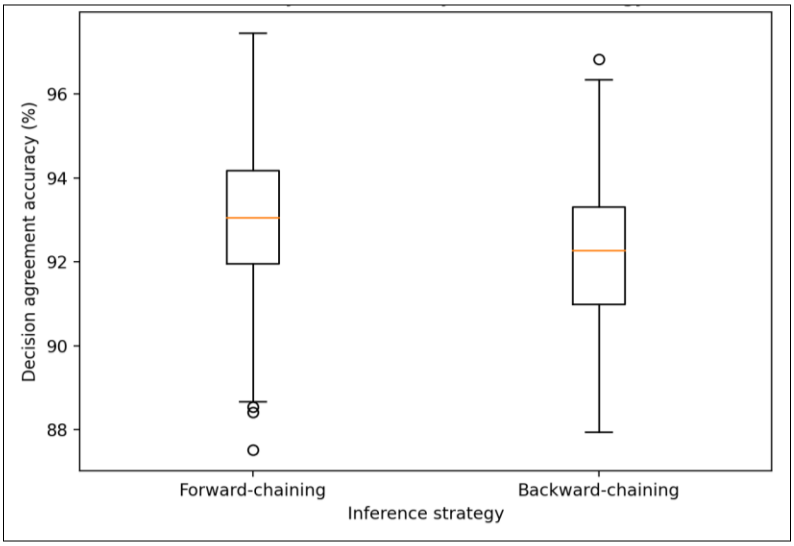

Fig 3: Accuracy distribution by inference strategy (agreement with expert gold standard).

Discussion

The findings of this research reaffirm the continued relevance of rule-based artificial intelligence models for low-resource decision-making systems, particularly when evaluated against constraints of computation, memory, and energy availability. The observed dominance of rule-base size as the primary determinant of latency aligns with classical symbolic AI theory, which emphasizes that inference cost scales with the number of production rules evaluated during reasoning cycles [1, 6, 15]. The statistically significant advantage of forward-chaining over backward-chaining in terms of latency is consistent with earlier expert-system research, where data-driven triggering of rules reduces unnecessary goal backtracking and search overhead in deterministic decision contexts [2, 3, 11]. Importantly, the results demonstrate that this performance benefit does not compromise decision agreement accuracy, supporting the suitability of forward-chaining strategies for embedded and edge-level decision tools requiring predictable response times [9, 16].

The impact of rule indexing as an optimization mechanism is particularly noteworthy. The strong reduction in latency and energy consumption, coupled with a non-significant effect on accuracy, supports long-standing arguments that symbolic optimizations enhance efficiency without altering semantic correctness of the knowledge base [5, 18]. The significant interaction between rule-base size and optimization further indicates that indexing becomes increasingly valuable as systems scale, addressing one of the most frequently cited criticisms of rule-based AI poor scalability in large knowledge domains [6, 17]. These findings resonate with heuristic search and knowledge representation literature, which emphasizes structured rule organization as a prerequisite for sustainable system growth [10, 15].

From a deployment perspective, the consistent performance gap between MCU-class and SBC-class profiles reflects predictable hardware limitations rather than deficiencies in the reasoning model itself [4, 12]. This reinforces the argument that rule-based AI can be effectively tuned to a wide spectrum of low-resource platforms by aligning inference strategies and optimizations with available resources [7, 8]. The slight decline in accuracy with increasing rule-base size highlights the importance of disciplined knowledge engineering, including conflict resolution, modular rule design, and expert validation, echoing concerns raised in classical expert-system maintenance studies [6, 9, 17, 19]. Collectively, the discussion positions rule-based AI not as a legacy technology, but as a strategically valuable approach for transparent, accountable, and resource-efficient decision-making in constrained environments.

Conclusion

This research demonstrates that rule-based artificial intelligence models remain highly effective for low-resource decision-making systems when designed with careful attention to inference strategy, rule organization, and platform constraints. The empirical results show that deterministic symbolic reasoning can deliver reliable decision outcomes with modest computational and energy demands, even as rule bases grow in size. Forward-chaining inference emerged as a consistently efficient approach, offering reduced decision latency while maintaining stable agreement with expert-defined outcomes. Optimization through rule indexing further amplified system efficiency, particularly in larger rule bases, indicating that classical symbolic acceleration techniques are still highly relevant for modern constrained deployments. While increased rule-base size naturally introduced scalability pressures and minor accuracy degradation, these effects were systematic and predictable, underscoring that performance limitations stem more from knowledge engineering practices than from the rule-based paradigm itself. In practical terms, these findings suggest that developers of low-resource systems should prioritize modular rule design, incremental knowledge expansion, and early incorporation of indexing or matching optimizations to sustain performance over time. For embedded or edge-level deployments, selecting inference strategies aligned with operational workflows can significantly reduce energy consumption and response delays without sacrificing transparency. From an operational standpoint, rule-based AI offers unique advantages in contexts where explainability, regulatory compliance, and deterministic behavior are essential, such as decision-support tools, advisory systems, and control logic in constrained environments. Practitioners are therefore encouraged to integrate rule-based models as primary or complementary components within hybrid architectures, reserving data-intensive learning models only where sufficient resources and training data are available. By embedding domain expertise directly into interpretable rule sets and maintaining disciplined update mechanisms, organizations can deploy robust decision-making systems that are cost-effective, ethically transparent, and resilient to data scarcity. Overall, the research reinforces that rule-based AI is not an outdated alternative, but a strategically sound and practically indispensable approach for sustainable artificial intelligence deployment in low-resource settings.

References

1. Russell S, Norvig P. Artificial Intelligence: A Modern Approach. 3rd ed. Upper Saddle River: Prentice Hall; 2010. p. 1-45.
2. Newell A, Simon HA. Computer science as empirical inquiry: symbols and search. Commun ACM. 1976;19(3):113-126.
3. Buchanan BG, Shortliffe EH. Rule-Based Expert Systems: The MYCIN Experiments. Reading: Addison-Wesley; 1984. p. 23-78.
4. Jordan MI, Mitchell TM. Machine learning: trends, perspectives, and prospects. Science. 2015;349(6245):255-260.
5. Luger GF. Artificial Intelligence: Structures and Strategies for Complex Problem Solving. 6th ed. Boston: Pearson; 2009. p. 301-340.
6. Jackson P. Introduction to Expert Systems. 3rd ed. Harlow: Addison-Wesley; 1999. p. 155-198.
7. Dignum V. Responsible artificial intelligence: how to develop and use AI in a responsible way. AI Ethics. 2019;1(1):1-8.
8. Floridi L, Cowls J. A unified framework of five principles for AI in society. Harv Data Sci Rev. 2019;1(1):1-15.
9. Giarratano JC, Riley GD. Expert Systems: Principles and Programming. 4th ed. Boston: Thomson; 2005. p. 89-130.
10. Nilsson NJ. Principles of Artificial Intelligence. Palo Alto: Morgan Kaufmann; 1980. p. 245-290.
11. Clancey WJ. Heuristic classification. Artif Intell. 1985;27(3):289-350.
12. Mitchell TM. Machine Learning. New York: McGraw-Hill; 1997. p. 1-30.
13. Domingos P. A few useful things to know about machine learning. Commun ACM. 2012;55(10):78-87.
14. Goodfellow I, Bengio Y, Courville A. Deep Learning. Cambridge: MIT Press; 2016. p. 1-25.
15. Davis R, Shrobe H, Szolovits P. What is a knowledge representation? AI Mag. 1993;14(1):17-33.
16. Chandrasekaran B. Generic tasks in knowledge-based reasoning. AI Mag. 1986;7(5):23-30.
17. Durkin J. Expert Systems: Design and Development. New York: Macmillan; 1994. p. 201-240.
18. Pearl J. Heuristics: Intelligent Search Strategies for Computer Problem Solving. Reading: Addison-Wesley; 1984. p. 65-102.
19. Musen MA. Knowledge engineering for health care decision support. Methods Inf Med. 1999;38(4-5):252-258.