Evidence Alignment and Transfer Boundaries in Resource-Efficient V2X Perception And Reinforcement Learning For Software Reasoning

Abstract

Research spanning resource-efficient V2X perception and reinforcement learning for software reasoning increasingly joins methods that were developed for different objects and decisions. Here, motion-aware approximate temporal memory for energy-efficient neural perception is compared with curriculum-aware reinforcement learning over code-lineage graphs to determine which claims can travel across those boundaries and which remain context dependent. Two target papers are triangulated against 12 locally validated publications. The comparison follows temporal redundancy, token selection, quantization, communication latency, safety assurance and deliberately separates mechanistic interpretation from performance ranking, because the latter can conceal incompatible experimental or operational conditions. The synthesis shows that temporal redundancy cannot be interpreted independently of token selection, while quantization determines whether an apparent improvement remains meaningful outside the original setting. The strongest claims are therefore those that expose sensitivity, failure conditions, and residual uncertainty. The resulting framework supports reproducible comparison while preserving differences between study designs, and it identifies concrete points at which transfer claims should be narrowed or retested.

Keywords

Resource-Efficient V2X Perception And Reinforcement Learning For Software Reasoning; Temporal Redundancy; Token Selection; Quantization; Communication Latency; Safety Assurance

1. Introduction

Research on resource-efficient V2X perception and reinforcement learning for software reasoning sits at an interface between scientific ambition and practical constraint. The field seeks not only stronger nominal performance but also explanations that explain both success and failure. That challenge is especially visible in comparing motion-aware approximate temporal memory for energy-efficient neural perception with curriculum-aware reinforcement learning over code-lineage graphs through evidence alignment and transfer boundaries. A result that is compelling in one experimental medium, dataset, cohort, geography, or system load can lose force under a different data-generating process. Consequently, synthesis must preserve the unit of analysis and the decision context. It must also distinguish a mechanism from a correlation, a benchmark advantage from a deployment advantage, and an interpretable diagnostic from a causal explanation.

This review uses five analytical lenses: temporal redundancy, token selection, quantization, communication latency, safety assurance. The lenses were selected because they jointly cover what is designed, how it is measured, and what must remain stable for the conclusion to transfer. The organizing principle is representation, objective, and evaluation protocol. Under that principle, methodological novelty is necessary but insufficient; the comparison also audits the baseline, uncertainty treatment, and resource or hazard boundary. The article is a literature synthesis and reports neither a new empirical study nor invented measurements or metrics.

2. Methodological Foundations

A useful conceptual decomposition begins with the input evidence, the transformation applied to it, the output used for evaluation, and the decision that follows. In resource-efficient V2X perception and reinforcement learning for software reasoning, these stages are often optimized separately even though errors propagate across them. An expressive model can magnify noise or bias already present in its evidence; an apparently accurate output can still be poorly calibrated for the final decision. For that reason, ablation evidence should accompany aggregate performance. A strong comparison states what the protocol fixes and what it lets move, then selects metrics that correspond to the intended use rather than to convenience alone.

The second foundation is boundary awareness. Temporal redundancy defines the local design problem, while safety assurance determines whether the design can survive outside that boundary. Intermediate lenses such as token selection and quantization connect those ends by exposing trade-offs that a single score conceals. This is where negative results and controlled perturbations locate the useful boundary of an explanation. For applied work, documentation of deployment constraints is therefore part of the scientific result, not an administrative afterthought.

3. Comparative Evaluation

The target study DORM-01, MotiMem: Motion-Aware Approximate Memory for Energy-Efficient Neural Perception in Autonomous Vehicles [1], addresses a concrete part of resource-efficient V2X perception and reinforcement learning for software reasoning through motion-aware approximate temporal memory for energy-efficient neural perception. It identifies a representation, mechanism, or evaluation boundary needed to compare claims. Against the focus on comparing motion-aware approximate temporal memory for energy-efficient neural perception with curriculum-aware reinforcement learning over code-lineage graphs through evidence alignment and transfer boundaries, the paper is read as a case whose boundary must be retained: transfer may depend on its sensing regime, preparation, workload, population, or benchmark. The synthesis preserves that scope and tests its assumptions against neighboring work.

The target study YAA-03, Evo-CuRL: Curriculum-Aware Reinforcement Learning over Code Lineage Graphs for Software Engineering Reasoning [2], addresses a concrete part of resource-efficient V2X perception and reinforcement learning for software reasoning through curriculum-aware reinforcement learning over code-lineage graphs. It identifies a representation, mechanism, or evaluation boundary needed to compare claims. Against the focus on comparing motion-aware approximate temporal memory for energy-efficient neural perception with curriculum-aware reinforcement learning over code-lineage graphs through evidence alignment and transfer boundaries, the paper is read as a case whose boundary must be retained: transfer may depend on its sensing regime, preparation, workload, population, or benchmark. The synthesis preserves that scope and tests its assumptions against neighboring work.

Across the focal studies, temporal redundancy should be interpreted jointly with token selection. A design can improve one yet redistribute uncertainty rather than remove it. The practical comparison is a matrix whose rows are operating conditions and whose columns are evidence quality, computation or process cost, robustness, and consequence of failure. Such a matrix prevents an attractive mechanism from being detached from the circumstances that support it. It also clarifies which ablations are informative: an ablation should challenge a mechanistic claim, not simply produce a score.

Quantization provides the bridge from method to decision. At minimum, evaluation should distinguish nominal behavior from stress response and show dispersion alongside averages, and test plausible alternative explanations. Where distribution shift is central, calibration or sensitivity is as important as ranking. Where costs or hazards are asymmetric, utility-weighted assessment is more revealing than undifferentiated accuracy. These choices do not make studies identical; they make the reasons for their differences auditable.

4. Architecture and Learning Objectives

A first comparison set connects PerceptNet-V2X: Perception Network for Vehicle to Everything Scenarios in Autonomous Driving [3]; Reinforcement Learning Model Based on Regret for Multi-Agent Conflict Games [4]; Extensible Heterogeneous Collaborative Perception in Autonomous Vehicles with Codebook Compression [5]. These publications were retrieved and verified through DOI or publisher metadata, and their titles directly intersect the problem boundary of resource-efficient V2X perception and reinforcement learning for software reasoning. Together they illuminate temporal redundancy: some emphasize representations or mechanisms, whereas others foreground validation and application conditions. The useful comparison is not a leaderboard across incompatible settings, but a structured reading of what each study controls, leaves variable, and treats as a meaningful outcome. Divergence can then reveal a change in scale, data process, endpoint, or operating constraint.

The neighboring evidence base brings together Stabilizing independent multi-agent reinforcement learning via curriculum-based iterative self-play [6]; Research on Cross-modal and Semantic Collaborative Unsupervised Domain Adaptation for All-weather Autono... [7]; Multi-hop temporal knowledge graph reasoning with multi-agent reinforcement learning [8]. These publications were retrieved and verified through DOI or publisher metadata, and their titles directly intersect the problem boundary of resource-efficient V2X perception and reinforcement learning for software reasoning. Together they illuminate token selection: some emphasize representations or mechanisms, whereas others foreground validation and application conditions. The useful comparison is not a leaderboard across incompatible settings, but a structured reading of what each study controls, leaves variable, and treats as a meaningful outcome. Divergence can then reveal a change in scale, data process, endpoint, or operating constraint.

A complementary strand is visible in Autonomous Aerial Vehicle Object Detection Based on Spatial Perception and Multiscale Semantic and Detai... [9]; Abmarl: Connecting Agent-Based Simulations with Multi-Agent Reinforcement Learning [10]; V2X-Sim: Multi-Agent Collaborative Perception Dataset and Benchmark for Autonomous Driving [11]. These publications were retrieved and verified through DOI or publisher metadata, and their titles directly intersect the problem boundary of resource-efficient V2X perception and reinforcement learning for software reasoning. Together they illuminate quantization: some emphasize representations or mechanisms, whereas others foreground validation and application conditions. The useful comparison is not a leaderboard across incompatible settings, but a structured reading of what each study controls, leaves variable, and treats as a meaningful outcome. Divergence can then reveal a change in scale, data process, endpoint, or operating constraint.

For triangulation, the literature includes Curriculum-Learning-Guided Multi-Agent Deep Reinforcement Learning for N-1 Static Security Prevention an... [12]; Intelligent Road Management System for Emergency Autonomous Buses (Eabs) Along with Emergency Autonomous... [13]; Curriculum-assisted multi-agent reinforcement learning for scalable V2X resource allocation [14]. These publications were retrieved and verified through DOI or publisher metadata, and their titles directly intersect the problem boundary of resource-efficient V2X perception and reinforcement learning for software reasoning. Together they illuminate communication latency: some emphasize representations or mechanisms, whereas others foreground validation and application conditions. The useful comparison is not a leaderboard across incompatible settings, but a structured reading of what each study controls, leaves variable, and treats as a meaningful outcome. Divergence can then reveal a change in scale, data process, endpoint, or operating constraint.

5. Deployment and Governance

For implementation, the review suggests a staged workflow. First, define the decision and its failure cost. Second, specify the evidence and operating envelope. Third, choose a method whose inductive assumptions match that evidence. Fourth, test communication latency and safety assurance under controlled perturbations. Finally, retain provenance so that an unexpected outcome can be traced to data, configuration, material batch, environmental condition, or user interaction. This sequence keeps comparing motion-aware approximate temporal memory for energy-efficient neural perception with curriculum-aware reinforcement learning over code-lineage graphs through evidence alignment and transfer boundaries connected to observable requirements.

The broader implication is that responsible progress in resource-efficient V2X perception and reinforcement learning for software reasoning depends on interfaces as much as components. Interfaces determine how uncertainty moves between measurement and inference, between algorithm and operator, or between microscopic design and system behavior. Joint work benefits from advance specification of acceptance criteria and falsifying evidence. When those agreements are explicit, optimization need not trade away validity, safety, or intelligibility without notice.

6. Limitations and Future Directions

This synthesis is limited by heterogeneity in terminology, study design, and reporting granularity. Metadata confirmation secures the citation trail but does not reproduce measurements or results. In addition, fast-moving work may change venue or version. Focal strings verified through Scholar were kept verbatim; uncertain fields were flagged in a parallel record.

Future work should preregister comparison protocols where feasible, publish machine-readable settings and test transfer under a substantively important shift in deployment constraints. Research should also group adverse cases by process and consequence. For resource-efficient V2X perception and reinforcement learning for software reasoning, a valuable shared benchmark would combine baseline effectiveness, resource consumption, calibration, and disturbance response. Such a benchmark would reward designs that remain useful when evidence is incomplete or conditions change.

7. Conclusion

The literature on resource-efficient V2X perception and reinforcement learning for software reasoning is best understood as a connected evidence system rather than a collection of isolated scores. The focal studies show how comparing motion-aware approximate temporal memory for energy-efficient neural perception with curriculum-aware reinforcement learning over code-lineage graphs through evidence alignment and transfer boundaries; the additional literature reveals the assumptions required to interpret those contributions. Across temporal redundancy, token selection, quantization, communication latency, and safety assurance, the most durable principle is alignment: the representation or mechanism must match the problem, the metric must serve the decision while deployment remains inside the tested boundary. Progress built on that alignment is easier to reproduce, safer to transfer, and more informative when it fails.

References

1. Que, H., Liu, M., Xie, J., Gao, H., Sun, J., Xu, H., ... & Qiao, F. (2026). MotiMem: Motion-Aware Approximate Memory for Energy-Efficient Neural Perception in Autonomous Vehicles. arXiv preprint arXiv:2603.27108.

2. Tan, W., Li, Y., & Liang, W. (2026, June). Evo-CuRL: Curriculum-Aware Reinforcement Learning over Code Lineage Graphs for Software Engineering Reasoning. In Proceedings of the 2026 International Conference on Multimedia Retrieval (pp. 1327-1335).

3. Costa, B. T., Pereira, C., & Maykol Pinto, A. (2025). PerceptNet-V2X: Perception Network for Vehicle to Everything Scenarios in Autonomous Driving. IEEE Access, 13, 182645-182660. https://doi.org/10.1109/access.2025.3624285

4. XIAO, Z., & ZHANG, S. Y. (2009). Reinforcement Learning Model Based on Regret for Multi-Agent Conflict Games. Journal of Software, 19(11), 2957-2967. https://doi.org/10.3724/sp.j.1001.2008.02957

5. Soorchaei, B. E., Raftari, A., & Fallah, Y. P. (2025). Extensible Heterogeneous Collaborative Perception in Autonomous Vehicles with Codebook Compression. Robotics, 14(12), 186. https://doi.org/10.3390/robotics14120186

6. Akgün, O. (2026). Stabilizing independent multi-agent reinforcement learning via curriculum-based iterative self-play. Neurocomputing, 704, 134819. https://doi.org/10.1016/j.neucom.2026.134819

7. Du, Z. (2026). Research on Cross-modal and Semantic Collaborative Unsupervised Domain Adaptation for All-weather Autonomous Driving Perception Enhancement. Exploring Science Academic Conference Series, 14, 400-406. https://doi.org/10.70267/ic-aimees.20260400406

8. Bai, L., Chen, M., & Xiao, Q. (2024). Multi-hop temporal knowledge graph reasoning with multi-agent reinforcement learning. Applied Soft Computing, 160, 111727. https://doi.org/10.1016/j.asoc.2024.111727

9. Rao, W., Chen, S., & Li, D. (2025). Autonomous Aerial Vehicle Object Detection Based on Spatial Perception and Multiscale Semantic and Detail Feature Fusion. IEEE Access, 13, 42897-42909. https://doi.org/10.1109/access.2025.3547825

10. Rusu, E., & Glatt, R. (2021). Abmarl: Connecting Agent-Based Simulations with Multi-Agent Reinforcement Learning. Journal of Open Source Software, 6(64), 3424. https://doi.org/10.21105/joss.03424

11. Li, Y., Ma, D., An, Z., Wang, Z., Zhong, Y., Chen, S., et al. (2022). V2X-Sim: Multi-Agent Collaborative Perception Dataset and Benchmark for Autonomous Driving. IEEE Robotics and Automation Letters, 7(4), 10914-10921. https://doi.org/10.1109/lra.2022.3192802

12. Zhang, X., Li, Z., Quan, X., Cheng, K., & Yu, Y. (2026). Curriculum-Learning-Guided Multi-Agent Deep Reinforcement Learning for N-1 Static Security Prevention and Control. Energy Engineering, 123(9), 1-10. https://doi.org/10.32604/ee.2025.073912

13. bin naeem, a., & Jing, Y. (2022). Intelligent Road Management System for Emergency Autonomous Buses (Eabs) Along with Emergency Autonomous Cars (Eacs) Under Vehicle-to-Everything (V2x) Communication. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.4281509

14. Morshed, M., & Zaman Chowdhury, M. (2026). Curriculum-assisted multi-agent reinforcement learning for scalable V2X resource allocation. Physical Communication, 76, 103062. https://doi.org/10.1016/j.phycom.2026.103062