Communication-Aware Memory and Semantic Compression for Collaborative Autonomous Perception

Abstract

Resource-Efficient V2X Perception has become a test of how well researchers can connect performance with evidence quality, resource limits, and transfer across settings. The review develops an evidence-centered account of jointly managing temporal memory and semantic tokens under energy and bandwidth constraints. The review triangulates 3 focal papers with 11 independently retrieved publications validated against DOI-registration metadata. The analysis is organized around temporal redundancy, token selection, quantization, communication latency, and safety assurance. A central precaution is not to treat results from heterogeneous studies as exchangeable, the review compares problem definitions, methodological assumptions, and validation boundaries. Across the literature, the literature consistently implies that advances in resource-efficient V2X perception become credible when representation, objective, and evaluation protocol are evaluated together and when uncertainty about distribution shift is reported explicitly. The resulting framework links method selection to decision consequence and recurring validity threats, and proposes a research agenda centered on well-specified controls, robustness tests, and reproducible workflows.

Keywords

Resource-Efficient V2X Perception; Temporal Redundancy; Token Selection; Quantization; Communication Latency; Safety Assurance

1. Introduction

Scholarship concerning resource-efficient V2X perception links scientific ambition and practical constraint. The objective includes not only stronger nominal performance but also explanations of the conditions and mechanisms behind success. The issue is particularly acute for jointly managing temporal memory and semantic tokens under energy and bandwidth constraints. A mechanism demonstrated for a bounded material, dataset, population, scale, or operating trace can erode beyond the conditions that produced it. The review process consequently has to 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.

Five lenses organize the comparison: temporal redundancy, token selection, quantization, communication latency, safety assurance. These dimensions matter because they jointly cover what is designed, how it is measured, and what must remain stable for the conclusion to transfer. Its governing principle is representation, objective, and evaluation protocol. Under that principle, originality needs evidence beyond a headline metric; the comparison also evaluates comparator fairness, uncertainty reporting, and real operating constraints. The analysis is literature based and is not presented as a new experiment and supplies no synthetic performance claim.

2. Methodological Foundations

The evidence pathway can be divided into the input evidence, the transformation applied to it, the output used for evaluation, and the decision that follows. In resource-efficient V2X perception, these stages are often optimized separately even though errors propagate across them. Upstream noise may grow as it passes through a flexible model; an apparently accurate output can still be poorly calibrated for the final decision. Accordingly, ablation evidence should accompany aggregate performance. A useful evaluation makes explicit which factors remain invariant and which may change, then selects metrics that correspond to the intended use rather than to convenience alone.

The next requirement is control of scope. Temporal redundancy defines the local design problem, while safety assurance determines whether the design can survive outside that boundary. Trade-offs become visible through token selection and quantization connect those ends by exposing trade-offs that a single score conceals. Unexpected outcomes and perturbation analysis reveals the interval in which an explanation can still be defended. 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 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 jointly managing temporal memory and semantic tokens under energy and bandwidth constraints, the paper serves as a design case with explicit limits: 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 DORM-03, DinoLink: A Token-Centric Representation Compression Framework for Bandwidth-Constrained Collaborative V2X Perception [2], addresses a concrete part of resource-efficient V2X perception through saliency selection and token-centric compression for constrained V2X communication. It identifies a representation, mechanism, or evaluation boundary needed to compare claims. Against the focus on jointly managing temporal memory and semantic tokens under energy and bandwidth constraints, the paper serves as a design case with explicit limits: 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 while imposing a less visible trade-off on the other. The analysis can be summarized in a compact 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 further reveals which ablations are informative: ablation design should target causal attribution instead of numerical proliferation.

Quantization links technical design with operational choice. At the least, assessment should separate familiar inputs from perturbations and retain distributional summaries, 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 comparison choices do not require identical designs; they make the reasons for their differences auditable.

4. Architecture and Learning Objectives

For triangulation, the literature includes PerceptNet-V2X: Perception Network for Vehicle to Everything Scenarios in Autonomous Driving [3]; Extensible Heterogeneous Collaborative Perception in Autonomous Vehicles with Codebook Compression [4]; Scholarship concerning Cross-modal and Semantic Collaborative Unsupervised Domain Adaptation for All-weather Autono... [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. 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 design space is broadened by Autonomous Aerial Vehicle Object Detection Based on Spatial Perception and Multiscale Semantic and Detai... [6]; V2X-Sim: Multi-Agent Collaborative Perception Dataset and Benchmark for Autonomous Driving [7]; Intelligent Road Management System for Emergency Autonomous Buses (Eabs) Along with Emergency Autonomous... [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. 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 first comparison set connects A systematic literature review on ground vehicle and unmanned aerial vehicle collaborative routing and s... [9]; A Method for Mapping V2X Communication Requirements to Highly Automated and Autonomous Vehicle Functions [10]; CoDAC: Autonomous Obstacle Avoidance Optimization for Unmanned Surface Vehicle Clusters via Multi-Modal... [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. 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.

The neighboring evidence base brings together RSPMP: real-time semantic perception and motion planning for autonomous navigation of unmanned ground ve... [12]; Edge Intelligence for V2X Communications: Advances in Collaborative Perception and Privacy Preservation [13]. These publications were retrieved and verified through DOI or publisher metadata, and their titles directly intersect the problem boundary of resource-efficient V2X perception. 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 practice, five ordered decisions are useful. 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. The staged design keeps jointly managing temporal memory and semantic tokens under energy and bandwidth constraints connected to observable requirements.

The broader implication is that responsible progress in resource-efficient V2X perception is governed by interfaces as well as components. Interfaces determine how uncertainty moves between measurement and inference, between algorithm and operator, or between microscopic design and system behavior. Teams spanning disciplines should predefine acceptance rules and triggers for revising conclusions. When those agreements are explicit, efficiency can be improved without concealing losses in validity, safety, or interpretability.

6. Limitations and Future Directions

Several limitations qualify the synthesis, including heterogeneity in terminology, study design, and reporting granularity. Verified authorship and venue do not substitute for independent empirical replication. Publication status can change after metadata collection. The supplied focal strings remain preserved while questionable normalization is shown only as a candidate.

Further investigation ought to preregister comparison protocols where feasible, retain machine-readable setup details while testing at least one nontrivial shift in deployment constraints. Research should also connect failure frequency to an explanatory taxonomy. For resource-efficient V2X perception, a valuable shared benchmark would combine ordinary performance, operational burden, uncertainty calibration, and resilience. Such a benchmark would reward designs that remain useful when evidence is incomplete or conditions change.

7. Conclusion

The source base for resource-efficient V2X perception is better interpreted through the relationships among studies rather than a collection of isolated scores. The focal studies show how jointly managing temporal memory and semantic tokens under energy and bandwidth constraints; 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, assessment must correspond to the decision and deployment claims to the evaluated envelope. 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. Zhu, T., Que, H., Yao, H., Xu, H., & Bao, Z. (2026). DinoLink: A Token-Centric Representation Compression Framework for Bandwidth-Constrained Collaborative V2X Perception. arXiv preprint arXiv:2606.26398.
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. 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

5. 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

6. 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

7. 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

8. 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

9. Cai, L., Yang, Z., Wang, Q., Hu, M., & Lodewijks, G. (2026). A systematic literature review on ground vehicle and unmanned aerial vehicle collaborative routing and scheduling: collaborative mode, uncertainty, and future direction. Autonomous Transportation Research. https://doi.org/10.1016/j.atres.2026.08.001

10. Takacs, A., & Haidegger, T. (2024). A Method for Mapping V2X Communication Requirements to Highly Automated and Autonomous Vehicle Functions. Future Internet, 16(4), 108. https://doi.org/10.3390/fi16040108

11. Li, Y., Ye, J., Gao, L., & Cai, M. (2025). CoDAC: Autonomous Obstacle Avoidance Optimization for Unmanned Surface Vehicle Clusters via Multi-Modal Dynamic Perception and Collaborative Community Detection. IEEE Access, 13, 134552-134569. https://doi.org/10.1109/access.2025.3593236

12. Chen, D., Zhuang, M., Zhong, X., Wu, W., & Liu, Q. (2022). RSPMP: real-time semantic perception and motion planning for autonomous navigation of unmanned ground vehicle in off-road environments. Applied Intelligence. https://doi.org/10.1007/s10489-022-03283-z

13. Wang, S. (2025). Edge Intelligence for V2X Communications: Advances in Collaborative Perception and Privacy Preservation. Journal of Big Data and Computing, 3(4), 73-81. https://doi.org/10.62517/jbdc.202501409