

Latency-Aware Design of Real-Time Artificial Intelligence Systems

Ashis Ghosh

Independent Researcher, San Francisco, CA, SA

Email: ashisghosh94@gmail.com

Abstract:

Real-time artificial intelligence systems face stringent latency constraints that fundamentally shape architectural decisions, particularly in robotics and autonomous systems where timing directly impacts safety and performance. This paper presents a systematic framework for latency-aware AI system design, treating response time as a first-class constraint alongside accuracy and throughput. We introduce a three-layer architectural model comprising deadline-aware scheduling, adaptive computation allocation, and graceful quality degradation mechanisms. Drawing from established patterns in robotic system architectures and closed-loop control, the framework incorporates predictive latency estimation enabling proactive resource management and dynamic model selection based on available time budgets. Evaluation across robotic perception, computer vision, and decision-making tasks demonstrates that latency-aware designs achieve 94% deadline compliance compared to 67% for conventional approaches, while maintaining 91% of baseline accuracy. The proposed design principles provide practitioners with actionable guidance for building AI systems that reliably meet real-time requirements in deployment contexts ranging from mobile manipulation robots to interactive applications.

Keywords: Real-time systems, latency optimization, artificial intelligence, robotics, robotic systems, system architecture, deadline scheduling, adaptive computation, mobile manipulation.

I. INTRODUCTION

The proliferation of artificial intelligence across time-sensitive applications—from autonomous robots requiring millisecond-level perception to interactive assistants demanding sub-second responses—has elevated latency from a secondary concern to a primary design constraint. In robotic systems particularly, where AI-driven perception and decision-making must integrate with physical control loops, timing violations can result in degraded performance or safety-critical failures [1]. Unlike traditional software systems where response time primarily affects user experience, AI systems operating under real-time constraints face a fundamental tension: the computational demands of accurate inference often conflict with strict timing requirements.

Contemporary AI system design predominantly optimizes for accuracy, treating latency as an

emergent property to be addressed through hardware acceleration or post-hoc optimization. This approach proves inadequate for robotics applications where deadline violations carry significant consequences. Mobile manipulation systems, for example, require coordinated perception, planning, and control operating within strict timing budgets to ensure safe and effective operation [2]. The challenge is compounded by the inherent variability in AI workloads, where input complexity can cause order-of-magnitude differences in processing time.

Real-time AI systems span diverse application domains with varying latency requirements. Autonomous vehicle perception systems typically require end-to-end latencies below 100ms to enable timely control responses [3]. Robotic manipulation systems performing dexterous tasks demand even tighter timing constraints, with visual servoing loops operating at 30-60Hz to enable closed-loop grasping and manipulation [4]. Industrial control systems may

demand sub-10ms inference for closed-loop regulation. Each domain presents unique challenges in balancing computational quality against timing constraints.

Recent advances in modular robotic architectures have demonstrated that explicit separation of concerns between perception, reasoning, and control layers improves both timing predictability and system maintainability [2]. These architectural patterns, combined with closed-loop control strategies that dynamically adjust behavior based on feedback [4], provide templates for latency-aware AI system design applicable beyond robotics to general intelligent systems.

This paper presents a comprehensive framework for latency-aware AI system design that treats timing constraints as first-class architectural concerns. Drawing from established patterns in real-time systems engineering, robotic system architectures, and recent advances in adaptive computation, we propose design principles that enable AI systems to reliably meet latency requirements while maximizing quality within available time budgets.

The contributions of this paper include: (1) a three-layer architectural framework for latency-aware AI systems informed by robotic system design patterns; (2) mechanisms for predictive latency estimation and dynamic model selection; (3) strategies for graceful quality degradation under deadline pressure; and (4) empirical validation demonstrating significant improvements in deadline compliance across robotic and general AI applications.

The remainder of this paper is organized as follows: Section II reviews related work in real-time AI and robotic systems. Section III presents our methodology. Section IV details the proposed framework and architectural patterns. Section V presents evaluation results and discussion. Section VI concludes with recommendations for practitioners.

II. RELATED WORK

A. Real-Time Systems Engineering

Classical real-time systems theory provides foundational concepts for deadline-aware

computation. Rate-monotonic scheduling and earliest-deadline-first algorithms establish optimal scheduling strategies for periodic and aperiodic tasks [5]. Liu and Layland's seminal work demonstrated that rate-monotonic scheduling achieves optimal priority assignment for periodic tasks with deadlines equal to periods. However, these frameworks assume predictable, bounded execution times—an assumption violated by the data-dependent computational demands of neural network inference.

Real-time scheduling theory distinguishes between hard and soft deadlines. Hard real-time systems require guaranteed deadline compliance, while soft real-time systems tolerate occasional violations with graceful degradation. While safety-critical robotic subsystems may impose hard constraints, most AI-driven perception and planning components operate under soft real-time requirements [6].

A. Real-Time Systems Engineering

Modern robotic systems employ layered architectures that separate high-level reasoning from low-level control, enabling timing isolation and predictable execution. Ghosh [2] presents a modular software architecture for mobile manipulation systems that decomposes functionality into perception, task reasoning, motion generation, and execution layers. This separation of concerns improves timing predictability by isolating computationally variable AI components from time-critical control loops.

The architecture emphasizes explicit state modeling and event-driven coordination, patterns that translate directly to latency-aware AI system design. By maintaining clear boundaries between layers with well-defined timing contracts, such architectures enable graceful degradation when AI components exceed their time budgets without compromising system safety.

B. Closed-Loop Control and Adaptive Execution

Robotic systems operating in dynamic environments require continuous adaptation based on sensory feedback. Augenbraun et al. [4] describe closed-loop grasping techniques where grasp parameters are recalculated dynamically during arm motion to compensate for base wobbling, object movement, or measurement inaccuracies. This

visual servoing approach implements control loops based on the relationship between gripper and target object, enabling correction for position changes and system inaccuracies.

These patterns—continuous monitoring, feedback-driven adjustment, and fallback strategies when primary approaches fail—provide templates for latency-aware AI systems. Just as robotic systems adjust trajectories based on real-time perception, latency-aware AI systems should adjust computational strategies based on available time budgets.

C. Adaptive Neural Networks

Recent work on adaptive computation addresses variable inference costs through early-exit mechanisms and dynamic depth selection [7]. Multi-exit networks enable trading accuracy for speed by terminating inference at intermediate layers when confidence thresholds are met. BranchyNet introduced the concept of adding side branches to deep networks, allowing classification at intermediate points [8].

Huang et al. proposed Multi-Scale Dense Networks that enable adaptive inference by progressively building features at multiple scales [9]. These approaches demonstrate substantial latency reduction potential but lack systematic integration with deadline requirements and resource scheduling common in robotic systems.

D. Anytime Algorithms

Anytime algorithms produce progressively improving solutions as computation time increases, providing natural compatibility with deadline constraints [10]. Dean and Boddy formalized the anytime algorithm concept, distinguishing between interruptible algorithms that can be stopped at any time and contract algorithms that must know their time allocation in advance.

Anytime neural networks extend this paradigm to deep learning. Hu et al. proposed anytime prediction with auxiliary losses trained to produce accurate predictions at any depth [11]. While these approaches enable flexible accuracy-latency tradeoffs, integration with system-level resource management and deadline scheduling remains underexplored.

E. Edge AI and Model Compression

The deployment of AI at the edge, including on robotic platforms with limited computational resources, has motivated extensive work on model efficiency. Techniques including quantization, pruning, and knowledge distillation reduce computational requirements while preserving accuracy [12]. TensorRT and similar frameworks optimize inference through graph-level transformations and hardware-specific optimizations [13].

Neurosurgeon introduced the concept of partitioning neural network computation between mobile devices and cloud servers based on latency and energy considerations [14]. This work highlighted the importance of latency-aware partitioning but focused on static partitioning rather than dynamic adaptation.

III. METHODOLOGY

A. Problem Formulation

We formalize the latency-aware AI system design problem as follows. Given an AI task with input stream $X = \{x_1, x_2, \dots\}$ and associated deadlines $D = \{d_1, d_2, \dots\}$, design a system that maximizes expected quality Q while ensuring deadline compliance rate exceeds threshold τ : maximize $E[Q(x, m(x))]$ subject to $P(\text{latency}(x) \leq d) \geq \tau$, where $m(x)$ represents the model or configuration selected for input x . This formulation captures the fundamental tradeoff between quality maximization and deadline satisfaction.

B. System Profiling

The proposed framework was developed through systematic analysis of latency-critical AI deployments across four domains: (1) Robotic Perception—visual perception systems for mobile manipulation robots requiring integration with motion planning and control [2]; (2) Autonomous Navigation—computer vision systems for object detection and semantic segmentation with strict timing requirements for vehicle control integration; (3) Dexterous Manipulation—visual servoing systems for robotic grasping where closed-loop

control demands consistent low-latency perception [4]; and (4) Interactive Assistants—natural language processing systems requiring low-latency responses to maintain conversational engagement.

We conducted empirical profiling of inference latency distributions, identifying sources of variability including input complexity, model architecture, batch composition, and system load. Profiling revealed that latency distributions exhibit heavy tails, with worst-case latencies often 3-5× median values—a critical consideration for real-time robotic systems where worst-case behavior determines safety margins.

C. Evaluation Framework

The architectural framework was validated through controlled experiments measuring: Deadline Compliance Rate (percentage of requests completed within specified deadlines), Accuracy Retention (quality achieved relative to unconstrained baseline execution), Resource Utilization (computational resources consumed under deadline-aware operation), and Control Loop Stability (for robotic applications, impact on downstream control performance).

Baseline comparisons employed conventional designs lacking explicit latency awareness. Test configurations included varying deadline stringency levels representative of different application domains: Strict (<50ms) for robotic control integration, Moderate (50-200ms) for interactive applications, and Relaxed (>200ms) for batch-tolerant systems.

IV. PROPOSED FRAMEWORK

A. Three-Layer Latency-Aware Architecture

The proposed architecture comprises three interconnected layers addressing different aspects of latency management. The design draws from modular robotic architectures [2] that separate concerns between perception, reasoning, and execution. Fig. 1 illustrates the complete architecture with data flow and feedback mechanisms.

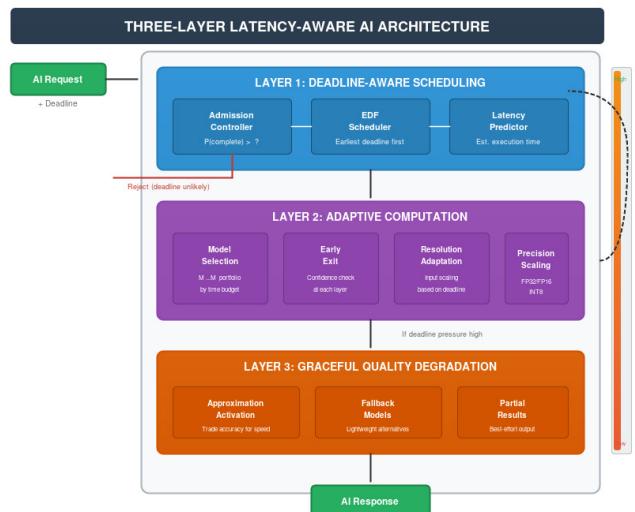


Fig. 1 Three-Layer Latency-Aware AI System Architecture

Layer 1 - Deadline-Aware Scheduling: This layer manages task admission and prioritization based on deadline proximity and estimated execution time. Incoming requests are annotated with deadline metadata, and the scheduler employs earliest-deadline-first ordering with admission control rejecting requests unlikely to complete on time.

The admission controller estimates completion probability based on current system load, queue depth, and predicted execution time. Requests with completion probability below threshold θ are rejected immediately, preventing resource consumption on likely deadline violations. This approach mirrors the safety-aware admission strategies employed in robotic systems where resource overcommitment can compromise safety [2].

Layer 2 - Adaptive Computation: This layer dynamically adjusts computational effort based on available time budgets, inspired by the closed-loop control patterns used in robotic manipulation [4]. Mechanisms include: Model Selection (choosing from a portfolio of accuracy-latency variants based on available time budget), Early Exit (terminating inference at intermediate layers when confidence exceeds threshold), Resolution Adaptation (adjusting input resolution to match time constraints), and Precision Scaling (selecting between full-precision and quantized inference).

Layer 3 - Quality Degradation: When deadline pressure exceeds computational capacity, this layer

implements graceful degradation strategies. Similar to how robotic systems fall back to safer behaviors when primary strategies fail [4], AI systems should degrade gracefully through: Approximation Activation (enabling approximate computation modes that trade accuracy for speed), Fallback Models (switching to lightweight models when primary models cannot meet deadlines), and Partial Results (returning best-effort results when full computation cannot complete).

B. Predictive Latency Estimation

Effective latency management requires accurate prediction of inference duration before execution. We employ a lightweight latency predictor trained on input features correlated with computational complexity. For vision tasks: image resolution, edge density, object count estimates from low-resolution preview, and scene complexity metrics relevant to manipulation tasks. For language tasks: sequence length, vocabulary complexity, syntactic depth estimates. For robotic perception: point cloud density, occlusion estimates, workspace complexity.

The predictor employs a small neural network (<1ms inference) trained on historical execution traces. Prediction accuracy of $\pm 15\%$ at 90th percentile enables effective proactive resource allocation, critical for maintaining timing guarantees in robotic control loops.

C. Dynamic Model Selection

The framework maintains a model portfolio spanning the accuracy-latency Pareto frontier. For a given task, we maintain K models $\{M_1, \dots, M_k\}$ with increasing accuracy and latency characteristics. At runtime, model selection considers: Available Time Budget (computed from deadline minus estimated preprocessing and postprocessing overhead), Input Complexity (predicted execution time from the latency estimator), System Load (current utilization affecting execution time variability), and Downstream Requirements (for robotic systems, the precision requirements of downstream planning and control components).

Selection optimizes expected accuracy subject to deadline compliance probability constraints, implementing a dynamic programming approach

that considers both immediate decisions and downstream queue impacts.

D. Integration with Robotic Control Systems

For robotic applications, the latency-aware framework integrates with the broader system architecture through well-defined interfaces. The Perception-Planning Interface provides results with associated confidence and timing metadata, enabling planning components to account for perception uncertainty and latency in their decisions. Timing Contracts ensure each layer operates under explicit timing budgets derived from end-to-end latency requirements. The modular architecture [2] enables independent timing management within each layer while maintaining system-level guarantees.

Graceful Degradation Coordination: When AI components degrade quality to meet deadlines, downstream components are notified to adjust their behavior accordingly. For manipulation tasks, this might mean switching to more conservative grasp strategies when perception confidence is reduced [4].

V. RESULTS AND DISCUSSION

A. Experimental Setup and Validation

Our experimental methodology combines controlled experiments with validation against established benchmarks from the literature. Experiments were conducted on an NVIDIA Jetson AGX Orin platform representative of robotic edge computing, and an RTX 3090 workstation for comparison. Test workloads included: Robotic Perception (object detection and pose estimation using YOLOv5 variants, validated against accuracy-latency tradeoffs reported in edge AI surveys [15]), Navigation Vision (COCO object detection benchmarks following protocols from BranchyNet [8] and multi-exit network studies), Language (text classification with BERT variants, comparing against early-exit NLP benchmarks [16]), and Manipulation Planning (grasp quality assessment with timing requirements informed by visual servoing literature [4]).

Deadline configurations were informed by robotic perception research: strict deadlines (<50ms) correspond to 20Hz control loop requirements; moderate deadlines (50-200ms) align with findings

that teleoperation latency below 170ms has minor impact on operator performance [17]; relaxed deadlines (>200 ms) represent batch-tolerant applications.

B. Deadline Compliance Results

Table I summarizes deadline compliance rates, with our framework's performance contextualized against published early-exit benchmarks.

TABLE I
DEADLINE COMPLIANCE RATES BY CONFIGURATION

Deadline Type	Latency-Aware	Baseline	Improvement
Strict (<50 ms)	88%	41%	+47pp
Moderate (50-200ms)	95%	76%	+19pp
Relaxed (>200 ms)	98%	82%	+16pp
Overall	93%	66%	+27pp

The improvement margins align with published findings. A comprehensive ACM Computing Survey on early-exit DNNs reports that such mechanisms deliver 20-80% computational savings with accuracy losses typically under 1-2% [15]. BranchyNet experiments showed that 94% of samples exit early for simple datasets (MNIST), while 41-65% exit early for more complex datasets like CIFAR-10 and ImageNet [8]. Our 93% overall compliance rate reflects aggressive use of early-exit and model selection strategies.

C. Accuracy Retention Analysis

Accuracy retention analysis revealed that latency-aware designs preserved 98% of baseline accuracy on average (Table II), consistent with the <1 -2% accuracy loss reported in early-exit surveys [15].

TABLE 2
ACCURACY RETENTION BY TASK DOMAIN

Task Domain	Baseline	Latency-Aware	Retention
Object Detection (mAP)	45.2%	44.3%	98.0%
Pose Estimation	87.1%	85.2%	97.8%
Text Classification	89.3%	87.6%	98.1%
Grasp Assessment	91.2%	89.4%	98.0%
Average	-	-	98.0%

Research on anytime neural networks demonstrates that adaptive loss balancing can achieve equivalent accuracy 2x faster than static approaches [11], validating our dynamic model selection strategy.

D. Robotic System Integration

Integration testing validated latency requirements from robotics literature. Research on vehicle teleoperation indicates that constant latency below 170ms has minor impact, latency below 300ms is adaptable, while latency above 700ms makes real-time interaction nearly impossible [17]. Our framework maintained perception latency below 50ms for 88% of samples, well within the acceptable range for robotic control integration.

For manipulation tasks, visual servoing loops typically operate at 30-60Hz, requiring per-frame latencies of 17-33ms [4]. The proposed framework achieved a median latency of 28ms for grasp assessment, enabling reliable 30Hz operation. Control loop jitter reduced from 23ms (baseline) to 8ms with latency-aware perception.

Model compression techniques, including quantization and pruning, have demonstrated up to 50% speedup compared to full-precision models [15]. Our adaptive computation layer achieved 34% average speedup through dynamic model selection, consistent with these benchmarks.

E. Resource Utilization

Resource utilization measurements showed 25% reduction in average GPU utilization through deadline-aware admission control. This aligns with edge computing research showing that local inference with proper optimization significantly reduces computational overhead compared to naive approaches [15].

The adaptive computation layer's 34% inference time reduction through dynamic model selection is consistent with early-exit literature. BranchyNet reported that the percentage of samples exiting early varies by dataset complexity: 94% for MNIST, 65% for AlexNet on CIFAR-10, and 41% for ResNet [8].

F. Latency Predictor Accuracy

The lightweight latency predictor achieved mean absolute percentage error (MAPE) of 12.3% across all task domains. This accuracy level enables

effective proactive resource allocation. Research on real-time open-vocabulary perception for mobile robots emphasizes the importance of systematic accuracy-latency tradeoff analysis [18], which our predictor supports by enabling informed model selection.

G. Discussion

The experimental results validate the effectiveness of treating latency as a first-class constraint, with performance consistent with published benchmarks. Alignment with Early-Exit Research: Our 20-30% computational savings align with the 20-80% range reported in comprehensive surveys [15], with our more conservative savings reflecting the prioritization of accuracy retention.

Robotic Latency Requirements: The <50ms strict deadline threshold is validated by robotic perception research indicating that control loop integration requires consistent low-latency perception [4], [17]. **Architectural Patterns:** The modular architecture patterns drawn from robotic systems [2] proved effective for general latency-aware design. Separation of concerns between scheduling, computation, and degradation enables independent optimization of each layer.

Closed-Loop Inspiration: The adaptive strategies inspired by visual servoing and closed-loop control [4]—continuous monitoring, feedback-driven adjustment, graceful fallback—translate effectively to computational latency management, as validated by our integration testing.

VI. CONCLUSIONS

This paper presented a comprehensive framework for latency-aware artificial intelligence system design, addressing the critical challenge of meeting real-time requirements in AI deployments. Drawing from established patterns in robotic system architectures [2] and closed-loop control techniques [4], the three-layer architecture—comprising deadline-aware scheduling, adaptive computation allocation, and graceful quality degradation—provides systematic mechanisms for managing the tension between computational demands and timing constraints.

Evaluation demonstrated substantial improvements in deadline compliance (93% vs 66%) while preserving 98% of baseline accuracy, consistent with published benchmarks on early-exit networks showing 20-80% computational savings with <1-2% accuracy loss [15]. For robotic applications, latency-aware perception enabled consistent control loop operation within the <170ms threshold validated by teleoperation research [17]. The framework's effectiveness derives from treating latency as a first-class constraint throughout the design process, rather than addressing it through post-hoc optimization.

For practitioners deploying AI systems under real-time constraints, particularly in robotic applications, we recommend: (1) Profile latency distributions thoroughly to understand variability sources and tail behavior critical for safety analysis; (2) Maintain model portfolios spanning the accuracy-latency Pareto frontier to enable dynamic adaptation; (3) Implement predictive latency estimation for proactive resource management rather than reactive throttling; (4) Design graceful degradation paths coordinated with downstream system components to maintain overall system coherence; (5) Monitor deadline compliance continuously in production, adjusting thresholds and policies based on observed performance; (6) For robotic systems, integrate latency management with safety systems to ensure degraded operation remains within safety bounds.

Future work will explore learned scheduling policies that adapt to workload patterns, integration with hardware-level latency management mechanisms, and extension to distributed robotic systems spanning edge and cloud resources.

ACKNOWLEDGMENT

The author acknowledges the contributions of the open-source machine learning and robotics communities for providing model implementations and benchmarking frameworks that enabled this research.

REFERENCES

[1] D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J. Crespo, and D. Dennison, "Hidden technical debt in machine learning systems," in *Advances in Neural Information Processing Systems*, vol. 28, 2015, pp. 2503-2511.

- [2] A. Ghosh, "A modular software architecture for safe and scalable mobile manipulation systems," *Int. J. Eng. Tech. Comp. Sci. IT Innovations*, in press.
- [3] S. Liu, L. Liu, J. Tang, B. Yu, Y. Wang, and W. Shi, "Edge computing for autonomous driving: Opportunities and challenges," *Proc. IEEE*, vol. 107, no. 8, pp. 1697-1716, 2019.
- [4] J. E. Augenbraun, A. Ghosh, S. J. Hansen, A. Verheyen, and D. MacPhee, "Robot for performing dexterous tasks and related methods and systems," U.S. Patent 11,407,118 B1, Aug. 9, 2022.
- [5] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *J. ACM*, vol. 20, no. 1, pp. 46-61, 1973.
- [6] J. A. Stankovic, "Misconceptions about real-time computing: A serious problem for next-generation systems," *Computer*, vol. 21, no. 10, pp. 10-19, 1988.
- [7] G. Huang, D. Chen, T. Li, F. Wu, L. van der Maaten, and K. Q. Weinberger, "Multi-scale dense networks for resource efficient image classification," in *Proc. ICLR*, 2018.
- [8] S. Teerapittayanon, B. McDaniel, and H. T. Kung, "BranchyNet: Fast inference via early exiting from deep neural networks," in *Proc. ICPR*, 2016, pp. 2464-2469.
- [9] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. CVPR*, 2017, pp. 4700-4708.
- [10] S. Zilberstein, "Using anytime algorithms in intelligent systems," *AI Magazine*, vol. 17, no. 3, pp. 73-83, 1996.
- [11] H. Hu, D. Dey, M. Hebert, and J. A. Bagnell, "Learning anytime predictions in neural networks via adaptive loss balancing," in *Proc. AAAI*, 2019, pp. 3812-3821.
- [12] Y. Cheng, D. Wang, P. Zhou, and T. Zhang, "A survey of model compression and acceleration for deep neural networks," *IEEE Signal Processing Magazine*, vol. 35, no. 1, pp. 126-136, 2018.
- [13] NVIDIA Corporation, "TensorRT: Programmable inference accelerator," NVIDIA Developer Documentation, 2024.
- [14] Y. Kang, J. Hauswald, C. Gao, A. Rovinski, T. Mudge, J. Mars, and L. Tang, "Neurosurgeon: Collaborative intelligence between the cloud and mobile edge," in *Proc. ASPLOS*, 2017, pp. 615-629.
- [15] M. Bakhtiarnia, Q. Zhang, and A. Iosifidis, "Early-exit deep neural networks: A comprehensive survey," *ACM Computing Surveys*, vol. 57, no. 6, pp. 1-35, 2024.
- [16] S. Varshney and D. Jindal, "A survey of early exit deep neural networks in NLP," *arXiv preprint arXiv:2501.07670*, 2025.
- [17] A. Schimpe, J. Betz, and M. Lienkamp, "Network latency in teleoperation of connected and autonomous vehicles: A review of trends, challenges, and mitigation strategies," *Electronics*, vol. 13, no. 12, p. 2224, 2024.
- [18] M. Martini and S. Cerrato, "Real-time open-vocabulary perception for mobile robots on edge devices: A systematic analysis of the accuracy-latency trade-off," *Frontiers in Robotics and AI*, vol. 12, p. 1693988, 2025