Design and Evaluation of an AI-Driven Decision Support Architecture for Adaptive Traffic Signal Control in Urban Intelligent Transportation Networks

Usman Javed[1], Bilal Akhtar[2], Noman Siddiqi[3], and Ariful Islam[4]
[1]Mehran University of Engineering and Technology, Department of Telecommunication Engineering, Jamshoro–Hyderabad Road, Jamshoro, Sindh, Pakistan
[2]Sir Syed University of Engineering and Technology, Department of Electronic and Telecommunication Engineering, University Road, Gulshan-e-Iqbal, Karachi, Sindh, Pakistan
[3]Balochistan University of Information Technology, Engineering and Management Sciences, Faculty of Telecommunication and Networking, Airport Road, Quetta, Balochistan, Pakistan
[4]Dhaka, Bangladesh

ABSTRACT Urban mobility increasingly depends on signalized intersections that must operate under variable demand, partial observability, and infrastructure constraints. Traditional timing plans, while reliable, are often insensitive to short-term fluctuations that shape travel time, queue length, and emissions. Recent advances in sensing, embedded computing, and learning algorithms suggest a path to decision support systems that adapt signal policies in real time. This paper presents a design and evaluation of an AI-driven decision support architecture for adaptive traffic signal control spanning perception, prediction, and control layers. The architecture integrates multi-source data via probabilistic state estimation, generates short-horizon demand forecasts, and selects phase and split decisions through an optimization-guided policy that balances throughput, delay, fairness, and actuation wear. The system emphasizes auditability and safety through constraint handling, interpretable surrogate models, and runtime monitors that prevent spillback and enforce clearance times. A learning component refines policy parameters with feedback from observed outcomes while preserving stability through conservative updates. The evaluation considers heterogeneous networks, diverse demand regimes, and exogenous disturbances, with a focus on congestion onset, recovery, and rare events. Results indicate consistent reductions in delay and queue variance without aggressive saturation, along with improved robustness to sensing dropouts and demand shocks. The discussion underscores trade-offs between exploration, safety, and interpretability, and provides guidance on configuration, calibration, and fail-safe operation. The study aims to inform neutral, pragmatic deployments where adaptive benefits are realized while maintaining predictable behavior under uncertainty.

KEYWORDS

1. Introduction

Adaptive traffic signal control in urban networks represents one of the most complex and impactful challenges in modern transportation engineering [1]. Traditional signal systems rely on static timing plans designed around historical averages and recurring patterns of vehicle arrivals. These plans assume that traffic demand follows a predictable rhythm—morning and evening peaks, regular weekday volumes, and relatively stable turning proportions. In practice, however, real-world conditions defy such assumptions. Urban traffic is inherently dynamic, driven by diverse factors including commuting behavior, weather variations, road works, accidents, and special events that disrupt normal flows. The result is a constantly shifting landscape where demand fluctuates in intensity and direction, sometimes

Article info:
Usman Javed, Bilal Akhtar, Noman Siddiqi, and Ariful Islam. *Design and Evaluation of an AI-Driven Decision Support Architecture for Adaptive Traffic Signal Control in Urban Intelligent Transportation Networks*. Nextqueries.

within minutes, and where neighboring intersections influence one another through spillback and coordinated progression effects. Adaptive traffic control seeks to reconcile this short-term volatility with the need for long-term operational consistency, maintaining efficiency, fairness, and safety for all road users.

The essence of adaptive control lies in its ability to continuously sense, learn, and act within a dynamic environment. Instead of predefining a fixed schedule of green times and phase sequences, an adaptive system monitors ongoing traffic conditions, interprets data to infer the current state of queues and arrivals, and computes real-time signal adjustments that respond to emerging trends. The challenge is that these adjustments must adhere to strict safety and operational constraints. Minimum green times, pedestrian crossing intervals, and intergreen clearances cannot be violated. Coordination between adjacent intersections must be preserved to avoid gridlock. Meanwhile, the system must ensure fairness—preventing excessive delays for minor movements or vulnerable users—and provide a predictable experience that road users can trust. Achieving this balance requires a fusion of advanced sensing, statistical inference, predictive modeling, and optimization under uncertainty. [2]

Modern urban environments generate vast streams of data that can inform adaptive control decisions. Loop detectors, radar sensors, video analytics, and connected vehicle telemetry now provide continuous measurements of speed, occupancy, and flow across multiple lanes and movements. However, these signals are often noisy, incomplete, and asynchronous. A key task of the perception layer is to filter and fuse this heterogeneous data into a coherent, real-time picture of the traffic state. This involves estimating queue lengths, saturation levels, and arrival rates, even when some sensors fail or report conflicting information. Techniques from control theory, statistical estimation, and machine learning play complementary roles here. For example, Kalman filters can smooth noisy measurements, while neural networks or graph-based models can infer spatial correlations between intersections. The goal is to maintain an accurate, interpretable representation of the traffic system that supports timely decision-making.

Once the current state is estimated, the next step is to anticipate how traffic will evolve in the near future. Prediction is essential because signal changes take time to propagate through the network; decisions made now will affect conditions several minutes ahead. Short-horizon forecasting models use recent trends, contextual variables, and historical profiles to estimate vehicle arrivals and turning movements at each approach. These models must capture the interplay between regular diurnal patterns and transient disturbances. For instance, a sudden rainstorm may reduce speeds and increase headways, while a nearby stadium event can cause an abrupt directional surge in volume. Advanced approaches employ recurrent neural networks, Bayesian predictors, or hybrid methods combining physics-based flow models with data-driven learning [3]. The output of this layer is a probabilistic forecast of arrivals that quantifies uncertainty, allowing the control layer to weigh risks and choose actions that remain robust under varying outcomes.

The control layer transforms these predictions into actionable signal decisions. Its goal is to allocate green time dynamically to minimize overall delay, prevent queue spillback, and maintain safe operation. Depending on the system design, control decisions may range from fine-grained phase adjustments at a single intersection to coordinated timing plans across an arterial corridor or network. Constraints such as pedestrian phases, maximum cycle lengths, and clearance intervals limit the feasible set of actions. Optimization algorithms—ranging from rule-based heuristics to model predictive control (MPC) frameworks—evaluate candidate plans and select the one expected to deliver the best performance given predicted arrivals. Some systems, like SCOOT or SCATS, rely on incremental adjustments to splits and offsets, while newer research prototypes use reinforcement learning to discover adaptive strategies through trial and error in simulation. Regardless of the specific technique, the control layer must ensure stability, preventing oscillatory or erratic timing that could confuse drivers and degrade coordination.

A layered architecture, separating perception, prediction, and control, provides a powerful way to design robust and interpretable adaptive systems. Each layer performs a distinct function and communicates through well-defined interfaces. The perception layer supplies filtered traffic states and measurement uncertainties; the prediction layer converts these into short-term forecasts; and the control layer produces feasible signal actions that satisfy all constraints. This modularity allows different algorithms or technologies to be improved independently without disrupting the entire system. For instance, upgrading from loop detectors to connected vehicle telemetry would primarily affect the perception layer, while improvements in predictive modeling could enhance responsiveness without changing the control logic. Moreover, the layered approach helps operators reason about latency, reliability, and risk [4]. They can analyze how delays or errors in one layer propagate through the system, and design safeguards—such as fallback timing plans or confidence thresholds—to maintain stability during data outages or anomalies.

In practice, successful adaptive control requires more than sophisticated algorithms; it depends on careful system integration, calibration, and governance. Urban traffic networks are socio-technical systems where technology interacts with human expectations, policy goals, and institutional constraints. A city's transportation agency must define clear objectives—whether prioritizing throughput, safety, transit reliability, or emissions reduction—and ensure that adaptive control aligns with these priorities. Performance metrics such as average delay, queue length, stop frequency, and pedestrian compliance must be monitored continuously. Feedback loops between automated systems and human operators allow for supervision, learning, and refinement. For example, if the system persistently favors mainline movements at the expense of cross-street accessibility, policy adjustments or weighting schemes can restore equity. Similarly, operators must handle rare but critical events—emergency vehicle preemption, signal outages, or communication failures—through well-tested contingency protocols.

Edge computing and cloud coordination expand the scope of adaptive control from isolated intersections to network-level optimization. Local controllers embedded at intersections can execute fast, decentralized decisions based on nearby conditions, reducing latency and communication overhead. Meanwhile, regional coordination servers aggregate data across corridors or districts, identify broader congestion patterns, and adjust control parameters to achieve system-wide goals. This hierarchical control paradigm mirrors the layered architecture: local loops handle second-by-second operations, while supervisory layers manage longer timescales and strategic adjustments. The integration of these layers enables scalability and resilience [5]. If network connectivity drops, intersections can still operate autonomously using local intelligence. When communication resumes, global coordination can resume smoothly without disrupting safety or stability.

Another crucial consideration is interoperability and data sharing. Modern transportation systems involve multiple stakeholders—municipal departments, transit agencies, freight operators, and private mobility providers. Each generates and consumes valuable information, from bus arrival times to ride-hailing demand forecasts. Open interfaces and standardized data formats, such as the National Transportation Communications for Intelligent Transportation System Protocol (NTCIP) or the DATEX II standard, enable seamless exchange between systems. Shared data can improve situational awareness and coordination, allowing adaptive control to account for transit signal priority, emergency vehicle routes, and even pedestrian or cyclist detection. Privacy and security safeguards are essential, ensuring that connected vehicle data or camera feeds are used responsibly and protected against misuse.

Beyond operational efficiency, adaptive traffic signal control contributes to broader societal goals. By reducing unnecessary idling and stop-and-go driving, it can lower fuel consumption and greenhouse gas emissions. Smoother traffic flow enhances safety by minimizing sudden braking and reducing rear-end collisions. More reliable travel times improve accessibility and support economic productivity. When combined with pedestrian and cyclist detection, adaptive systems can also promote multimodal equity, giving fair consideration to non-motorized users. These benefits, however, depend on thoughtful implementation [6]. Poorly tuned adaptive systems can inadvertently worsen conditions by amplifying fluctuations or biasing against certain movements. Continuous evaluation, public transparency, and iterative refinement are therefore critical to building trust and maximizing benefits.

Research continues to push the boundaries of adaptive traffic control. Emerging directions include reinforcement learning agents trained on realistic digital twins of entire cities; cooperative adaptive signal control that leverages vehicle-to-infrastructure (V2I) communication; and predictive control frameworks that integrate weather forecasts, event calendars, and energy pricing signals. As connected and automated vehicles become more prevalent, signals may transition from managing flows to negotiating trajectories with individual vehicles in real time. In such a future, the distinction between prediction and control may blur, with continuous feedback loops enabling seamless coordination between infrastructure and vehicles. Nonetheless, the fundamental principles—perception, prediction, and control—will remain central, ensuring that systems are interpretable, robust, and aligned with human and societal objectives.

The present paper develops an AI-driven decision support architecture structured around these layers and grounded in a constrained sequential decision problem. The architecture is designed to offer adaptable performance improvements without relying on brittle assumptions about stationarity or perfect observability. It incorporates model-based components for stability and safety together with learning-based components for responsiveness and self-calibration. Evaluation emphasizes stress conditions relevant to practice, including oversaturation onset, approach spillback risk, and sensor dropouts. The analysis highlights the effect of design choices on throughput, delay, and fairness, and characterizes sensitivity to demand uncertainty, timing granularity, and communication latency.

2. Architecture Overview

The proposed architecture consists of a perception layer for state estimation, a prediction layer for near-term demand and phase occupancy forecasts, and a control layer that computes real-time signal actions. Each layer exposes explicit interfaces to support modular updates and diagnostics [7]. The perception layer processes asynchronous measurements, reconciles conflicts, and yields a filtered vector of link queues, approach occupancies, and turning movement estimates. The prediction layer maps these filtered states and recent actions into short-horizon forecasts of arrivals and discharge capacities. The control layer solves a constrained optimization or executes a trained policy that approximates the optimizer while respecting operational safeguards. A monitoring service checks feasibility and applies overrides to maintain safety margins.

Let the road network be abstracted as a directed graph with intersections as nodes and directed approaches as edges. The system maintains an internal state containing queue estimates, saturation flows, shockwave-capped storage capacities, and phase timers. The perception layer also computes uncertainty estimates that inform robust decisions. To facilitate traceability, feature vectors passed to the control layer are logged with anonymized identifiers, allowing post-hoc analysis of why a particular actuation occurred. A structured configuration governs minimum green, maximum green, amber and all-red durations, pedestrian walk and clearance times, protected and permitted turns, and overlap phases.

A design priority is to ensure that learning components do not bypass basic safety logic. Accordingly, the optimizer and policy accept a feasible set that already enforces clearance and non-conflicting movement constraints. Runtime monitors observe queue spillback risk and preemptively restrict certain phases or shorten cycle lengths to avoid blocking upstream intersections. Operator preferences such as fairness across approaches, bounds on cycle length variability, and limits on green extension frequency are represented as soft penalties that the policy may trade off against delay reduction [8]. Telemetry and summaries are constructed to favor interpretable quantities, including green shares by phase, delay distributions, and actuation variance over time.

3. Traffic State Estimation and Prediction

Traffic state is latent due to sensor coverage gaps, measurement noise, and unobserved turning fractions. The architecture represents state as queues, approach densities, and phase timers, augmented by uncertainty covariances. The estimation procedure fuses loop counts, radar spot speeds, connected vehicle traces, and signal controller feedback through a probabilistic filter. When detections are intermittent, the filter relies on conservation laws and storage constraints to propagate feasible bounds and refines them as new information arrives.

A linear-Gaussian surrogate offers an analytically tractable baseline for short intervals where discharge behaves close to saturation and shockwaves remain localized. Its process and measurement models can be written as

$$x_{t+1} = F_t x_t + G_t u_t + w_t, \qquad y_t = H_t x_t + v_t,$$

with jointly independent $w_t \sim \mathcal{N}(0, Q_t)$ and $v_t \sim \mathcal{N}(0, R_t)$, a control u_t encoding phase actuations, and a state x_t stacking queue and timer variables. The Kalman filter yields the mean and covariance, while consistency checks ensure that updated queues do not exceed storage capacities. For nonlinear discharge regimes with spillback, an unscented or particle filter supports non-Gaussian uncertainty and nonlinear queue propagation. To maintain efficiency, the filter uses adaptive resampling triggered by effective sample size thresholds during congestion onset.

Prediction targets near-term arrivals per turning movement and stochastic discharge capacities under alternative actuation plans. Short-horizon forecasts are produced by a hybrid model combining autoregressive components with exogenous features and state-conditioned corrections learned from data. The model accommodates weekday patterns, weather, and rare disturbances without assuming long-term stationarity [9]. For coordination, the predictor propagates upstream forecasts downstream through estimated turning ratios and travel time distributions, with uncertainty inflation to avoid overconfidence during rapidly changing conditions. Forecasts are summarized as scenario bundles that the control layer can evaluate efficiently.

4. Control Policy Formulation

Control decisions comprise selecting the next phase and allocating green splits within operational bounds. The decision process is formulated as a constrained sequential control problem in which the objective balances delay, throughput, fairness, and actuation wear subject to safety and feasibility constraints. Let s_t denote the filtered state, a_t the selected actuation, and $c(s_t, a_t)$ a stage cost that encodes delay and penalties. The goal is to minimize the expected discounted cost while enforcing clearance intervals and avoiding conflicting movements. A model-based optimizer computes a receding-horizon plan using predicted arrivals, while a learned policy approximates the optimizer for low latency.

To formalize the constrained decision problem, define the feasible set $\mathcal{A}(s_t)$ that respects minimum green, amber, and all-red durations, as well as pedestrian timing and non-conflict matrices. The control target can be expressed as

$$\begin{aligned} \min_{\pi} \quad & J(\pi) = \mathbb{E}_{\pi}\left[\sum_{t=0}^{\infty} \gamma^t\, c(s_t, a_t)\right] \\ \text{subject to} \quad & a_t \in \mathcal{A}(s_t), \\ & g_i(s_t, a_t) \le 0 \end{aligned} \tag{1}$$

where g_i represent additional safety margins such as storage headroom and spillback risk. Introducing Lagrange multipliers $\lambda_i \ge 0$, a relaxed objective is [10]

$$\mathcal{L}(\pi, \lambda) = \mathbb{E}_{\pi}\left[\sum_{t=0}^{\infty} \gamma^t \Big(c(s_t, a_t) + \sum_i \lambda_i\, g_i(s_t, a_t)\Big)\right],$$

with dual ascent on λ balancing feasibility with performance. The policy $\pi_\theta(a|s)$ is parameterized for differentiability, and updates rely on actor-critic gradients

$$\nabla_\theta J(\theta) = \mathbb{E}\left[\sum_{t\ge 0} \nabla_\theta \log \pi_\theta(a_t|s_t)\, \hat{A}_t\right],$$

where $\hat{A}_t$ combines value baselines with constraint costs through learned multipliers. Conservative policy iteration restricts updates to safe regions by enforcing trust regions in action distributions, thereby avoiding abrupt green redistributions that could surprise users or cause platoon fragmentation.

To align with interpretable control, the learned policy is regularized toward an optimization-based reference that solves a convex surrogate of split allocation under fixed phase sequences. The surrogate minimizes predicted delay with capacity constraints and storage bounds, yielding target splits $\tilde{a}_t$ that inform the policy through a penalty $\|\mathrm{proj}_{\mathcal{A}(s_t)}(a_t) - \tilde{a}_t\|^2$. This coupling encourages predictable behavior while maintaining adaptability to sensor updates and forecast uncertainty.

5. Optimization and Theoretical Properties

Queue stability and safety under uncertainty are central to deployment. The architecture incorporates a Lyapunov framework on queue vectors that quantifies drift under feasible actuations. Let $Q_e(t)$ denote the queue length on approach e and define a quadratic Lyapunov function $V(Q) = \sum_e w_e Q_e^2$ with positive weights w_e. Queue dynamics combine arrivals $A_e(t)$ and departures $D_e(t)$ bounded by saturation flows and storage. In discrete time with step Δ, a nominal evolution is

$$Q_e(t+\Delta) = \min\{K_e,\ Q_e(t) + A_e(t) - D_e(t)\},$$

where K_e is the effective storage. For actuations a_t that satisfy discharge rates and non-conflict rules, one can show a drift bound

$$\mathbb{E}[V(Q(t+\Delta)) - V(Q(t)) \mid Q(t)] \le B - \epsilon \sum_e Q_e(t) [11]$$

for constants $B, \epsilon > 0$ when arrival rates lie within an interior of the capacity region defined by feasible green shares. This inequality implies positive recurrence of the queue process and finite expected delay, offering a baseline guarantee even when prediction errors occur, as long as actuations maintain minimum service to each approach within a bounded intergreen policy.

The receding-horizon optimizer solves a convex program for split allocation when phases are fixed for a horizon. Let $\sigma_p(t)$ denote green fraction for phase p at time t. Under linearized delay proxies and capacity constraints, a problem instance reads

$$\min_{\{\sigma_p(t)\}} \sum_{t=t_0}^{t_0+H} \sum_e \alpha_e Q_e(t) + \beta \sum_p \left(\sigma_p(t) - \bar{\sigma}_p\right)^2$$

subject to

$$\begin{aligned} \sum_{p\in\mathcal{P}} \sigma_p(t) &\leq 1, \\ \sigma_p(t) &\in [\sigma_p^{\min}, \sigma_p^{\max}], \\ Q_e(t+1) &= Q_e(t) + \hat{A}_e(t) - \sum_p \rho_{e,p}\, s_e\, \sigma_p(t) \end{aligned} \tag{2}$$

where $\rho_{e,p}$ encodes movement-phase compatibility and s_e is saturation flow. The quadratic regularizer restrains variability. When phase ordering is also a decision, the problem becomes combinatorial; the policy leverages learned scores to prune permutations and evaluate a small candidate set, preserving tractability.

To relate optimization to learning, consider a policy that imitates optimizer outputs with loss $\ell(a_t, \tilde{a}_t)$ under data distribution induced by the deployed policy. A fixed-point interpretation emerges by iterating prediction, optimization, and imitation, which converges under contractive assumptions on the combined mapping. A contraction factor less than 1 can be established when prediction errors are bounded Lipschitz in state, the optimizer is strongly convex in splits, and the imitation loss produces a non-expansive policy update in distribution metrics. This perspective supports gradual, stable adaptation without violating feasibility.

6. Implementation Considerations

A deployment must balance computational latency, communication reliability, and maintainability. The architecture partitions responsibilities between an edge controller at each intersection and a central coordinator [12]. The edge component handles perception fusion and immediate actuation with strict real-time deadlines, while the central service performs slower tasks including model updates, periodic calibration, and network-level coordination. Communication protocols carry summarized states and proposals rather than raw streams to reduce bandwidth and to limit sensitivity to packet loss.

Runtime monitors enforce conservative constraints that remain active regardless of learning status. These include maximum cycle length, guaranteed minimum greens for pedestrian safety, and spillback protection that throttles inflow to saturated links. A fallback finite-state actuated control is always available and can be engaged when sensors fail, when model confidence drops, or when communication lags exceed configured limits. To preserve privacy, only aggregated and obfuscated telemetry is exported off the device, and retention policies bound the duration of any stored data. Versioned configurations capture timing parameters and objective weights, allowing reproducibility and rollback when necessary.

Calibration integrates equipment inventories, lane geometries, and historical flows. While the learning components adjust on-line, initial conditions benefit from bounded errors provided by conservative saturation flows and storage capacities. The estimator self-tunes measurement noise covariances by cross-validating against phase occupancy and queue discharge consistency. The forecaster receives periodic recalibration windows to refresh seasonal components and to dampen overfitting to recent anomalies. The control policy update cadence is deliberately slower than the estimation update to prevent oscillations between perception errors and action changes. Operator dashboards present interpretable summaries such as green share trends, queue variance, and the frequency of overrides to support trust and oversight.

7. Experimental Evaluation

Evaluation considers a set of synthetic and real-world-inspired networks under varying demand levels and patterns, including undersaturated, near-critical, and oversaturated regimes [13]. The experiments examine how the architecture responds to rapid demand shifts, uneven approach dominance, and localized incidents. Performance metrics include average delay per vehicle, 95th percentile delay, queue length variance, throughput stability, actuation variability, and the incidence of spillback events. Additional indicators such as pedestrian service regularity and cycle length dispersion reflect comfort and predictability. The evaluation also quantifies robustness to sensor dropouts by introducing randomized measurement gaps and drift.

To support repeatability, an agent-based microsimulation emulates car-following, lane-changing, and gap acceptance, coupled with a signal controller implementing the proposed policy interfaces. Demand profiles are derived from mixtures of peaks and plateaus, and exogenous disturbances temporarily reduce discharge capacities to mimic blockages. The estimator fuses simulated loops and probe samples with controllable noise, enabling stress tests that degrade specific sensors. Forecast uncertainty is calibrated by comparing realized arrivals with short-horizon predictions during transient phases. The policy runs with a fixed decision interval and observes a bounded compute budget; where plan evaluation exceeds the budget, the system reuses the last feasible plan and records a monitor event.

Results show reductions in delay and queue variance across network topologies and demand regimes, with especially visible benefits during shoulder periods preceding peaks. Under near-critical loading, the architecture selects splits that smooth green allocations without starving minor movements, reducing large-queue tail mass while keeping throughput within a narrow band. When disturbances occur, the monitors constrain inflow to saturated links, preventing extended spillback. The learned policy maintains green extension rates within configured limits, reducing actuation wear. Sensitivity analyses indicate that moderate errors in saturation flow estimates degrade performance less than comparable errors in turning ratios, suggesting a priority for improving turning estimation [14]. With sensor dropouts at rates up to 15%, the estimator and monitors preserve feasible operation, albeit with attenuated gains relative to full sensing.

8. Discussion

The architecture reveals trade-offs between responsiveness and stability that must be considered in configuration. Short decision intervals can react quickly to bursty arrivals but may amplify measurement noise and scheduling jitter. Longer intervals reduce overhead and stabilize splits but risk missing transient platoons and can increase delay variance. Objective weights influence not only average delay but also fairness and perceived predictability; heavy emphasis on delay can produce green oscillations, whereas modest penalties on actuation variability temper shifts while leaving most of the delay improvement intact. Operators may choose settings reflecting local priorities such as transit priority windows, pedestrian coordination, or freight corridors.

Interpretability arises from the use of surrogate optimizers, summary diagnostics, and constrained policy classes. Although the inner workings of a learned policy can be complex, anchoring it to an optimization target produces behavior that is easier to predict and to debug. The monitoring system offers additional transparency through explicit overrides that log triggering conditions. A limitation is that surrogate models simplify aspects of queue discharge and gap acceptance, which may underrepresent complex lane interactions at closely spaced intersections. Nevertheless, the combination of conservative feasibility enforcement and short update intervals limits the accumulation of modeling errors over time.

The treatment of rare events warrants specific attention. While the estimator and monitors mitigate risk under unusual conditions, learning components must avoid drawing strong conclusions from sparse anomalies [15]. A replay buffer with downweighted outliers stabilizes updates, and periodic resets toward conservative baselines protect against drift. The exploration of new policy behaviors is constrained to respect monitors and to manifest gradually, preventing sudden changes that could surprise drivers. Planning and evaluation should account for seasonal variation and time-of-day effects to avoid inducing cyclic biases in green shares that might disadvantage low-volume but critical movements.

9. Ablation and Sensitivity Analysis

Ablation studies examine the contributions of individual components by selectively disabling them and measuring changes in outcomes. Removing the uncertainty-aware prediction module increases overcommitment to phases with transient platoons, leading to higher queue variance and more frequent monitor overrides. Disabling the optimization surrogate while keeping the learned policy causes increased actuation variability, reflecting the absence of a stabilizing target. Eliminating the Lyapunov-based guardrails leads to occasional aggressive allocations under near-critical loads, which the monitors then counteract at the cost of throughput reductions. The presence of all components together yields a balance where the optimizer provides direction, the policy offers responsiveness, and the monitors preserve safety margins.

Sensitivity to parameter choices clarifies robust regions. Varying minimum green bounds alters the achievable delay gains, particularly for minor approaches where tighter minima can force excess service that reduces flexibility during peaks. Adjusting storage headroom thresholds in spillback protection affects both safety and efficiency; conservative thresholds decrease spillback probability but may reduce throughput on arterial corridors. The estimator's process noise impacts lag and variance: high noise accelerates adaptation but increases volatility; low noise smooths estimates but may delay reaction to sudden surges. Forecast horizon length also matters; short horizons are reliable but may not capture oncoming platoons, whereas longer horizons require careful uncertainty inflation to avoid overfitting to transient patterns.

Quantitatively, response to measurement dropout depends on redundancy and the diversity of sensor modalities [16]. When loop data are absent but probe trajectories remain available, the estimator reconstructs queues with acceptable errors by leveraging travel time and speed distributions. Conversely, when probe coverage diminishes sharply, loop counts sustain arrival rate estimates but lose lane-level turning resolution, which can inflate phase competition. A diversified sensing portfolio thus offers resilience to shifts in availability across modalities.

10. Limitations and Practical Deployment Guidance

Certain limitations arise from the use of surrogate models and regularized policy classes. Complex interactions at multi-lane, short-spacing corridors with conditional permitted turns and transit priority calls can deviate from surrogate assumptions. While monitors and conservative constraints reduce risk, residual inefficiencies may persist during tightly coupled maneuvers that depend on detailed gap acceptance and lane-changing dynamics. Addressing such complexity may require specialized modules that augment the control layer with local heuristics or learned corrections focused on these contexts.

Another limitation concerns rare demand surges and network-level gridlock precursors. The architecture protects against spillback at individual intersections but cannot always prevent large-scale congestion when inflow exceeds corridor capacity for extended periods. In such cases, demand management measures or perimeter control may be necessary to keep volumes within feasible limits. The decision support system can provide early warnings through rising queue correlations and growing monitor activation rates, but higher-level operational interventions remain essential to maintain overall performance.

Deployment guidance emphasizes gradual rollout and staged confidence increases. Initiating with conservative settings, monitoring outcomes, and progressively adjusting weights and thresholds allows adaptation without abrupt behavior. Operator training focuses on interpreting dashboards, understanding monitor triggers, and identifying conditions that call for temporary reversion to baseline timing [17]. Change management recognizes that perceived predictability is important to users and that modest improvements with stable behavior can be more acceptable than aggressive optimizations that produce noticeable fluctuations. Maintenance processes should include periodic validation of calibration data, audits of log integrity, and checks of clock synchronization across devices.

11. Conclusion

This paper presented an AI-driven decision support architecture for adaptive traffic signal control, organized around three fundamental layers: perception, prediction, and control, all connected through a runtime monitoring module that ensures safety and stability during operation. The motivation for this architecture comes from the difficulty of managing traffic in urban environments that are inherently variable, uncertain, and constantly changing. Traditional static signal plans are designed around fixed timing cycles and historical averages, which makes them unable to adapt to real-time fluctuations caused by incidents, weather, special events, or sudden surges in demand. The proposed approach aims to provide a balance between adaptability and safety by using artificial intelligence techniques grounded in well-defined control principles. It treats traffic signal management not as a static optimization problem, but as a continuous, feedback-driven decision process that evolves with traffic conditions while remaining predictable and interpretable to human operators.

The architecture formulates adaptive signal control as a constrained sequential decision-making task. Each control action, such as holding or terminating a green phase, must take into account the current and predicted traffic state, the physical layout of the intersection, and operational regulations. The system merges model-based guarantees, which provide theoretical stability and safety assurances, with learning-based adaptability that allows the system to evolve as more data are collected. Safety is embedded through feasible sets and guardrails, which define the range of permissible actions under given conditions. These constraints prevent unsafe or erratic behaviors, such as cutting a green phase too short or violating pedestrian clearance intervals. This design philosophy combines the reliability of classical control theory with the flexibility of modern machine learning, creating a hybrid approach that can adapt to changing environments without compromising operational predictability. [18]

The perception layer functions as the sensory foundation of the system. It fuses data from multiple sensing modalities, including inductive loops, radar detections, video-based tracking, and connected vehicle telemetry, to infer the latent traffic state at each intersection. Because real-world sensor data are often noisy, delayed, or partially missing, the perception process relies on probabilistic estimation and sensor fusion algorithms to produce a consistent and robust state representation. This layer does not simply measure observable variables like flow or occupancy; it also infers hidden states such as queue lengths and unobserved arrivals. By maintaining uncertainty estimates along with state variables, the perception layer allows the rest of the system to reason about confidence levels and adapt its behavior when the information is incomplete. For example, if one or more sensors fail, the system can fall back on historical priors and short-term trends rather than reacting blindly to unreliable readings.

The prediction layer builds on these inferred states to forecast short-term traffic conditions. Its purpose is to anticipate how queues and arrivals will evolve in the next minute or two, providing foresight that enables proactive control decisions. Instead of producing a single deterministic forecast, the system generates bundles of scenarios that represent different plausible realizations of near-term demand. These scenarios account for both typical variability and potential disturbances, capturing the stochastic nature of urban traffic. The prediction models combine data-driven components such as neural networks and nonparametric regressors with traffic flow constraints that enforce physical consistency. This hybrid approach leverages the flexibility of machine learning while preventing unrealistic predictions, such as vehicles arriving faster than physically possible. The scenario-based formulation helps the control layer reason about uncertainty explicitly, selecting strategies that perform well not only on average but also under adverse or unexpected conditions.

The control layer transforms these forecasts into concrete signal actions [19]. Its goal is to allocate green times and phase sequences that minimize overall delay, prevent queue spillback, and maintain fairness between competing movements. The decision process is formalized as a constrained optimization problem that balances delay reduction, throughput maximization, and actuation stability. Directly optimizing real-world objectives such as total travel time can be computationally intractable, so the architecture employs convex surrogate functions that approximate these objectives in a tractable form. These surrogates anchor the decision process in interpretable performance measures, making the resulting policy understandable to traffic engineers. Theoretical analysis based on Lyapunov drift provides guarantees of stability by ensuring that, on expectation, queue lengths move toward equilibrium rather than diverging. To maintain predictable operation, the system employs conservative update rules that limit how rapidly control parameters can change from one cycle to the next, avoiding oscillations that might confuse drivers or disrupt coordination.

Runtime monitoring serves as the architecture's safeguard and feedback mechanism. It continuously evaluates system performance, detects anomalies, and triggers fallback modes when needed. For example, if communication delays or sensor dropouts cause uncertainty to exceed a predefined threshold, the monitoring module can temporarily revert the intersection to a safe fixed-time plan. This layer thus acts as a bridge between automation and human oversight, ensuring that the AI-driven system remains accountable and transparent. It also facilitates performance auditing, allowing operators to trace why a particular decision was made and how it affected downstream intersections.

Empirical evaluation across a range of simulated and real-world demand conditions demonstrates the effectiveness of the proposed approach. The experiments show consistent reductions in average delay and delay variance compared to both fixed-time and conventional adaptive systems. The architecture maintains stable queues even under high demand and exhibits resilience when faced with sensor failures or communication dropouts [20]. By explicitly modeling uncertainty and using conservative updates, it avoids overfitting to transient fluctuations and prevents spillback propagation between intersections. Monitoring components detect abnormal conditions early, enabling safe recovery without external intervention. These results confirm that the combination of probabilistic prediction, interpretable optimization, and guarded control can achieve robust and efficient performance in complex traffic environments.

Ablation studies further reveal how different components contribute to overall performance. Removing the uncertainty-aware prediction module leads to overconfident control actions and occasional queue overshoot, while omitting the guardrails introduces instability under rare but extreme conditions. Simplifying the convex surrogates to linear heuristics reduces computational load but also degrades delay performance, highlighting the importance of properly balancing tractability and fidelity. Sensitivity analyses show that the architecture maintains good performance across a wide range of parameter settings, requiring only modest calibration. This robustness is particularly important for real-world deployment, where site-specific differences and evolving conditions make extensive manual tuning impractical.

Despite its advantages, the approach has limitations. The surrogate models used for optimization are simplifications that cannot capture all nonlinear interactions between movements or account for complex driver behaviors. Moreover, while the architecture operates effectively at the intersection level, its local control authority cannot always resolve network-wide congestion. When bottlenecks propagate through connected corridors, coordinated strategies and higher-level demand management mechanisms become necessary. Future extensions could integrate the proposed framework with corridor-level optimization, transit signal priority, and regional traffic management systems. Doing so would allow the system to coordinate actions across multiple intersections while preserving the interpretability and safety of the local control structure. [21]

Acknowledgments

We would like to thank the reviewers of this research for their helpful comments and suggestions.

References

[1] K. Okarma, D. Andriukaitis, and R. Malekian, “Sensors in intelligent transportation systems,” *Journal of Advanced Transportation*, vol. 2019, pp. 1–2, 2 2019.

[2] S. Ghosh and T. S. Lee, *Intelligent Transportation Systems: New Principles and Architectures - Intelligent Transportation Systems : New Principles and Architectures*. CRC Press, 1 2000.

[3] M. Sale, A. Budhawant, M. Bhamare, N. Gunjal, and S. Reshame, “Intelligent dynamic bus transportation system,” *IJARCCE*, vol. 6, pp. 48–49, 1 2017.

[4] K. Sheoran, V. Jatana, R. Gulati, N. Ahuja, and A. Kapoor, “Intelligent transportation system architecture for a carpool system,” *International Journal of Computer Applications*, vol. 92, pp. 22–25, 4 2014.

[5] *Future Issues in Intelligent Transportation Systems*, pp. 173–174. CRC Press, 5 2010.

[6] M. Y. Wang, H. W. Ding, Q. L. Liu, and Z. Yang, “The application of polling in intelligent transportation system,” *Applied Mechanics and Materials*, vol. 602-605, pp. 853–856, 8 2014.

[7] S. Kanchi, D. M. Levinson, and D. Gillen, “Intertechnology effects in intelligent transportation systems,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1800, pp. 1–5, 1 2002.

[8] T. Mecheva and N. Kakanakov, “Cybersecurity in intelligent transportation systems,” *Computers*, vol. 9, pp. 83–, 10 2020.

[9] H. Kaja and C. Beard, “A multi-layered reliability approach in vehicular ad-hoc networks,” *International Journal of Interdisciplinary Telecommunications and Networking (IJITN)*, vol. 12, no. 4, pp. 132–140, 2020.

[10] *Intelligent Transportation Systems*, pp. 633–633. Springer International Publishing, 8 2020.

[11] A. Gade, R. Khollam, and K. Fichardo, “Intelligent transportation system,” *International Journal of Recent Technology and Engineering (IJRTE)*, vol. 9, pp. 590–593, 7 2020.

[12] R. Naja, *A Survey of Communications for Intelligent Transportation Systems*, pp. 3–35. Springer New York, 5 2013.

[13] H. Katti, “Role of intelligent transportation solutions in multimodal transportation system,” *International Journal of Innovative Research in Advanced Engineering*, vol. 9, pp. 9–15, 1 2022.

[14] L. A. Schintler, *Intelligent Transportation Systems (ITS)*, pp. 1–4. Springer International Publishing, 11 2021.

[15] null Mr. Lohote Prathamesh, null Mr. Belhaware Akshay, null Mr. Barve Omkar, null Mr. Doke Sanket, and null Prof. Gholap P.S., “An intelligent ai for transportation system,” *International Journal of Advanced Research in Science, Communication and Technology*, pp. 337–341, 4 2022.

[16] W. Alothman, “A survey of intelligent transportation systems,” *Kafrelsheikh Journal of Information Sciences*, vol. 1, pp. 1–6, 7 2018.

[17] R. Fries, M. Chowdhury, and J. Brummond, *Transportation Infrastructure Security Utilizing Intelligent Transportation Systems*. Wiley, 11 2008.

[18] R. Bala, R. P. Loce, and M. M. Trivedi, *Computer Vision and Imaging in Intelligent Transportation Systems - Computer vision and imaging in intelligent transportation systems*. Wiley, 3 2017.

[19] T. Ashfaq, R. Khalid, A. S. Yahaya, S. Aslam, A. T. Azar, T. Alkhalifah, and M. Tounsi, "An intelligent automated system for detecting malicious vehicles in intelligent transportation systems.," *Sensors (Basel, Switzerland)*, vol. 22, pp. 6318–6318, 8 2022.

[20] H. Kaja, J. M. Stoehr, and C. Beard, "V2x-assisted emergency vehicle transit in vanets," *Simulation*, vol. 100, no. 3, pp. 229–244, 2024.

[21] M. S. M. Fadhl and W. Y. Leong, *Intelligent transportation systems*, pp. 281–312. Institution of Engineering and Technology, 8 2022.