

Event-Triggered Tube Model Predictive Control for Robust Vehicle Motion Tracking Under Uncertain Dynamics and External Disturbances

Mathan Breane

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.

nathangreene@uab.edu

Ashwin Gill

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.

ashwin.gill333@ku.edu

Vishal M. Joshi

Department of Computer Science, University of North Texas, Denton, TX, USA.

vjoshi@unt.edu

Abstract

The deployment of autonomous and semi-autonomous ground vehicles places exacting demands on motion tracking systems that must operate in uncertain environments, under changing vehicle dynamics, and in the presence of external disturbances such as crosswinds, road surface irregularities, and sensor degradation. This paper presents a system-level analysis of event-triggered tube model predictive control as an architectural framework for robust vehicle motion tracking. Rather than focusing exclusively on algorithmic formulations, the discussion addresses the coupling between predictive control structures, disturbance compensation, computational resource allocation, communication infrastructure, and safety governance. The event-triggered mechanism is examined as a means of reducing computational and communicative load while preserving robustness through tube-based uncertainty propagation. The paper further considers structural trade-offs involving prediction horizon length, triggering thresholds, state constraint tightening, and real-time solver performance. Infrastructure implications for embedded deployment, failure handling, and lifecycle sustainability are analyzed alongside cross-domain comparisons with networked control and large-scale data-driven systems. The analysis shows that event-triggered tube model predictive control offers a principled path toward balancing robust constraint satisfaction with practical resource limitations. However, its successful deployment depends on governance mechanisms, validation procedures, and lifecycle management that extend beyond the control law itself. The paper contributes a synthetic interdisciplinary perspective intended to inform both engineering design and institutional decision-making for robust autonomous vehicle control.

Keywords

event-triggered control, tube model predictive control, robust vehicle motion tracking, uncertain dynamics, external disturbances, real-time control systems, safety assurance, infrastructure governance.

1. Introduction

Model predictive control has become a central approach for constrained vehicle motion control because it can explicitly handle multivariable dynamics, actuator limits, and safety constraints within a single optimization-based framework. Early stability and optimality results established the conceptual foundation for constrained predictive control and demonstrated that receding horizon strategies can preserve closed-loop stability under appropriate terminal conditions [1]. Subsequent developments in theory and computation expanded these ideas into practical design procedures for constrained systems [2]. Within this broader field, robust model predictive control has received particular attention because embedded vehicle platforms rarely possess exact models and are constantly exposed to uncertain operating conditions.

A powerful robust formulation is tube-based model predictive control, in which an ancillary controller keeps the actual state within a bounded neighborhood of a nominal trajectory, while the nominal trajectory is optimized subject to tightened constraints [3]. This approach provides a systematic separation between nominal planning and disturbance rejection, enabling rigorous constraint satisfaction despite bounded uncertainties. However, solving a predictive optimization problem at every time step can impose significant computational and communication burdens on vehicle platforms with limited embedded processing capability and shared network resources. Event-triggered control offers a complementary mechanism for reducing the frequency of control updates by executing tasks only when certain conditions indicate that a new update is necessary [4]. Introductory surveys of event-triggered and self-triggered control have clarified the benefits and limitations of aperiodic task execution in networked systems [5].

Vehicle motion tracking raises additional challenges because the controlled process is highly safety-critical, strongly nonlinear, and subject to rapidly changing external conditions. Predictive active steering schemes have shown that predictive control can manage vehicle stability in demanding maneuvers [7]. Tube-based nonlinear predictive approaches have further demonstrated the ability to maintain robust performance in semiautonomous ground vehicles under model mismatch and road disturbances [8]. At the same time, cross-domain data-intensive methods have shown that large-scale operational data can improve prediction and adaptation across different environments, suggesting that future vehicle control systems may benefit from hybrid strategies that combine physical modeling with learned environmental models [6]. The present paper integrates these threads by examining event-triggered tube model predictive control from a systems perspective.

2. System Context and Related Work

The larger context for robust vehicle motion tracking includes motion planning, perception, localization, and low-level actuation, all of which operate under real-time deadlines and shared computational resources. A comprehensive survey of motion planning and control for self-driving urban vehicles highlighted the interdependence between trajectory generation, tracking, and constraint management [9]. Within this architecture, robust tracking controllers must respond not only to nominal planned trajectories but also to unexpected deviations caused by vehicle loads, tire wear, road curvature, and aerodynamic disturbances. Reachability analysis has been used to verify the safe behavior of automated road vehicles by computing sets of states that can be reached under uncertain inputs and disturbances, thereby providing a complementary safety certification mechanism [10].

From a computer-controlled systems viewpoint, the implementation of control algorithms is inseparable from sampling, quantization, scheduling, and communication delays [11]. A predictive control algorithm that performs well in simulation may fail on an embedded platform if its computational demands exceed the available execution window. Predictive control with constraints remains a key engineering reference for understanding the practical requirements of constrained optimization in real-time settings [12]. More recent work has emphasized automatic generation of high-speed solvers and efficient interior point methods for embedded predictive control, enabling richer optimization problems to be solved on automotive-grade processors [13], [14]. Robust predictive formulations must also account for persistent disturbances that cannot be assumed to vanish over time. Early work on predictive control with restricted constraints under persistent disturbances established important principles for constraint tightening and feasibility retention [15].

Learning-based predictive control has emerged as a promising direction for improving vehicle models and disturbance predictions, but it also introduces new challenges in safety certification and transient performance. Safe learning frameworks attempt to combine data-driven model improvement with rigorous safety constraints, often by maintaining conservative bounds until confidence in learned components increases [16]. Control barrier functions provide another formal tool for ensuring that state trajectories remain within safe sets, and they have been used alongside predictive methods in safety-critical applications [17]. These developments form the intellectual backdrop for event-triggered tube model predictive control, which seeks to combine robust constraint satisfaction, computational efficiency, and formal safety guarantees.

3. Event-Triggered Tube Model Predictive Control Architecture

The architecture considered in this paper combines three interacting layers: a nominal model predictive controller, a disturbance rejection mechanism based on tube geometry, and an event-triggered scheduling layer. The nominal predictive controller solves a constrained optimization problem over a finite receding horizon using an approximate model of the vehicle. The tube layer constructs a set of permissible deviations around the nominal trajectory, reflecting the expected influence of model uncertainty and bounded external disturbances. The event-triggering layer monitors the difference between the actual vehicle state and the nominal predicted state, and it initiates a new optimization only when this difference approaches the boundary of the tube or when a triggering condition indicates that continued use of the previous plan would risk constraint violation.

This layered architecture has several structural advantages. First, it separates nominal trajectory generation from uncertainty handling, which simplifies both design and verification. Designers can select a nominal model that captures the dominant vehicle dynamics without attempting to model every possible disturbance source in precise detail. The tube can then be expanded or contracted according to the estimated uncertainty level. Second, it reduces the effective computational load because the optimization is not necessarily performed at every sampling instant. The event-triggering rule implements a form of resource-aware control in which computation becomes demand-driven rather than time-driven. Third, the architecture supports graceful degradation: if the event-triggered communication channel becomes congested or unavailable, the tube-based ancillary controller can continue to stabilize the vehicle around the last nominal trajectory for a limited period, provided that the tube has not been exited.

From a systems engineering standpoint, the event-triggered tube model predictive control architecture also introduces important design parameters that must be co-selected rather than independently optimized. The prediction horizon determines how far into the future the nominal controller plans, influencing both computational cost and robustness. The triggering threshold determines how frequently new predictions are requested, thereby shaping the trade-off between communication overhead and tracking accuracy. The tube width controls how much uncertainty is accommodated before the nominal plan must be revised. These parameters are coupled because a narrow tube may trigger frequent replanning, while a wide tube may tolerate larger tracking errors but also require more conservative constraint tightening. System-level design therefore requires an integrated analysis of control performance, computational scheduling, and communication bandwidth.

4. Robustness and Disturbance Rejection from a Systems Perspective

Robustness in vehicle motion tracking cannot be reduced to a single numerical margin. It involves the ability to maintain safe operation under bounded model mismatch, sensor noise, actuator degradation, and environmental disturbances. Tube-based methods provide a conceptual language for discussing robustness because they treat uncertainty as a bounded set rather than as a purely statistical quantity. The vehicle is not expected to follow an exact trajectory but rather to remain within a tube that is propagated forward in time. This perspective aligns well with safety requirements, because safety can be expressed as avoiding certain regions of the state and input space under all admissible disturbance realizations.

External disturbances such as crosswinds, potholes, and uneven road surfaces have both transient and persistent components. Persistent disturbances are particularly challenging because they may bias the vehicle away from the nominal trajectory even when the controller remains active. The event-triggered tube predictive control framework addressed in a recent formulation provides input-to-state stability guarantees for vehicle motion control under disturbances, supporting the claim that the tracking error remains bounded when the disturbance is bounded [18]. Such stability properties are essential for certifying that the system does not exhibit unbounded drift or oscillatory behavior when facing sustained environmental loads.

Reachability analysis complements the tube-based perspective by computing boundary sets that characterize the worst-case influence of disturbances over a finite horizon. Ellipsoidal techniques for reachability provide computationally tractable outer approximations of reachable sets, which can be used to validate whether the tube configuration is consistent with the actual uncertainty model [19]. When reachable sets remain inside the corresponding tubes, the nominal plan retains validity. When they approach the tube boundary, the event trigger should request a replanning operation. This relationship shows that event triggering, tube geometry, and reachability analysis are not independent tools but rather parts of a unified robustness architecture.

The formal notion of input-to-state stability offers a useful bridge between control theory and systems-level resilience. It expresses the idea that the future state of the system remains bounded relative to the magnitude of disturbances, with a transient influence of initial conditions that decays over time [20]. In an event-triggered tube predictive control setting, the triggering condition must be designed so that input-to-state stability is preserved even when control updates occur sporadically rather than periodically. The fact that such guarantees can be established indicates that event-triggered operation does not inevitably sacrifice robustness,

provided that the triggering logic is carefully coordinated with tube size and disturbance bounds [18].

5. Structural Trade-offs in Computation, Communication, and Control

The introduction of event-triggering into tube model predictive control transforms the traditional trade-off space. In time-triggered implementations, the computational workload is highly predictable but potentially excessive. The processor performs the same optimization at every sampling instant, even when the current plan remains valid and the disturbance level is low. In event-triggered implementations, the workload becomes data-dependent and bursty. A period of calm conditions may require very few optimization calls, while a sudden disturbance or rapid maneuver may produce a cluster of triggering events. This irregularity has important consequences for embedded system design, including scheduling policies, memory usage, and thermal management.

Communication presents a similar trade-off. In a distributed vehicle architecture, the predictive controller may reside on a central computing unit while sensor and actuator data traverse an in-vehicle network. Event-triggered operation can reduce the number of messages exchanged between sensing, computation, and actuation nodes, which is valuable in bandwidth-limited automotive networks. However, aperiodic communication can introduce jitter and make end-to-end latency harder to bound. Designers must ensure that the worst-case latency between a triggering event and the corresponding actuator update remains compatible with the tube size and the rate of disturbance growth.

The choice of optimization solver also interacts with event triggering. High-speed embedded solvers can reduce the computational penalty of replanning, thereby making more frequent triggering acceptable [13]. Efficient interior point methods for multistage predictive problems can handle the structured optimization problems that arise in receding horizon control [14]. However, solver efficiency is not uniform across all problem instances, and event triggering may sometimes concentrate difficult problem instances in a short time window. The control architecture must therefore include safeguards against computational overrun, such as warm-start procedures, early termination criteria, or fallback strategies based on the last feasible nominal plan.

The priority assigned to robustness, communication reduction, and computational effort can differ significantly across deployment contexts. A high-speed highway vehicle may require tight tracking and frequent replanning, while a low-speed logistics platform may tolerate larger tracking error in exchange for longer battery life and reduced processing load. The event-triggered tube predictive control architecture supports this context-specific adaptation through parameter selection. Nevertheless, it is difficult to guarantee that a single fixed parameterization will perform well across all operational modes. Adaptive triggering policies that adjust thresholds according to road curvature, vehicle speed, and environmental conditions are therefore an important area of ongoing investigation.

6. Infrastructure, Deployment, and Real-Time Implementation

Practical deployment of event-triggered tube model predictive control requires more than an algorithm. It demands an embedded software infrastructure capable of managing real-time tasks, sensor fusion, communication interfaces, and safety monitors. The control law must coexist with perception pipelines that may themselves operate at variable rates, with localization systems that provide intermittent corrections, and with actuator controllers that have their own internal dynamics. The overall system must be designed so that the predictive

controller receives sufficiently fresh state estimates and can deliver commands within deterministic deadlines [11].

The hardware platform strongly influences the achievable control performance. Automotive-grade processors vary widely in their numerical capability, memory hierarchy, and support for parallel computation. Model predictive control has historically been deployed on automotive systems with restricted computational resources, but advances in embedded optimization have made richer control formulations feasible [9], [12]. Event-triggered operation can reduce average utilization, but peak utilization remains a critical concern. In safety-critical systems, the platform must be provisioned for worst-case computational demand rather than average demand, which may reduce the apparent benefits of event triggering if the triggering policy is not carefully designed.

Deployment also involves integration with the vehicle communication architecture. Modern vehicles use multiple communication buses and service-oriented middleware to connect sensors, computing nodes, and actuators. Event-triggered control can reduce traffic on these buses, but it may also introduce irregular communication patterns that complicate network scheduling. Time-sensitive networking, priority-based arbitration, and redundancy mechanisms must be evaluated together with event-triggering policies. The control architecture must remain functional when message delays vary, when packets are lost, or when sensor data become temporarily unavailable. Tube-based robustness provides a natural buffer against these communication anomalies, but only if the tube width and triggering thresholds account for the worst-case network behavior.

The real-time implementation of robust predictive control also requires a clear separation between nominal plan generation and fallback behavior. When a new nominal plan is not available at a triggering instant, the system must continue operating using the previous plan and the ancillary controller. This fallback behavior can be made safe if the tube remains valid and the vehicle has not approached constraint boundaries. However, prolonged fallback operation may lead to divergence if disturbances continue to push the vehicle away from the stale nominal trajectory. The deployment architecture should therefore include multiple layers of safety response, including conservative degradation of speed, trajectory widening, and ultimately a safe stopping maneuver.

7. Fairness, Safety, and Governance Implications

Safety is the primary governance requirement for vehicle motion control, but fairness and accountability also enter the system design conversation. A robust motion tracking algorithm should not perform well only in well-mapped urban environments while degrading disproportionately in rural, poorly maintained, or heterogeneous road settings. Cross-domain learning and transfer methods can reduce performance disparities by enabling models trained in data-rich environments to be adapted to data-scarce ones [6]. Such approaches may also introduce concerns about data representativeness, bias in training distributions, and hidden dependencies on environmental features.

The safety case for event-triggered tube model predictive control must address both nominal performance and degraded modes. Regulatory bodies and insurance institutions increasingly expect autonomous vehicle developers to provide evidence that control algorithms satisfy safety requirements under defined disturbance assumptions. The tube-based framework supports this expectation by making uncertainty bounds explicit and by providing a mechanism for constraint tightening. Event triggering, however, complicates verification

because the system no longer executes control updates at fixed intervals. Certification procedures must therefore account for the possibility that a triggering event may be delayed or missed and must demonstrate that the resulting state remains within acceptable bounds.

Governance mechanisms should also define the responsibilities of different organizational actors in the development and deployment process. The control algorithm provider must document the assumptions on vehicle dynamics, disturbance bounds, sensor quality, and computational resources. The vehicle integrator must ensure that the deployed platform satisfies those assumptions. The fleet operator must monitor in-service performance and update uncertainty models as conditions change. An event-triggered tube predictive control system is not a static artifact but a dynamic socio-technical system whose performance depends on maintenance, software updates, and environmental monitoring. Clear governance structures are therefore necessary to ensure that robustness guarantees remain meaningful over the vehicle lifecycle.

Transparency is another governance concern. Decision-makers, auditors, and possibly vehicle occupants need to understand why the controller initiated a replanning operation or why it continued to use an older plan. Event-triggering rules can be recorded and time-stamped to create an audit trail of control decisions. This traceability can support post-incident analysis and ongoing safety assessment. It can also help identify systematic weaknesses in the triggering policy, such as excessive triggering in particular road segments or insufficient triggering during certain disturbance patterns. The integration of predictive control with data-driven environmental prediction [6] must be accompanied by data governance practices that protect privacy, ensure representativeness, and prevent inappropriate reuse of operational data.

8. Sustainability and Lifecycle Considerations

Sustainability in vehicle control systems extends beyond tailpipe emissions to include computational energy consumption, hardware longevity, and the environmental footprint of data acquisition and model development. Event-triggered tube model predictive control can contribute to energy efficiency by reducing the average computational load on embedded processors and lowering communication activity. Because the controller performs fewer optimization calls during routine driving, the processor may spend more time in low-power states, reducing thermal stress and potentially extending hardware life. These benefits depend on the triggering policy and the workload characteristics of the specific driving environment.

The lifecycle perspective also encompasses software maintenance and model updates. Vehicle dynamics change over time due to tire wear, suspension aging, and vehicle loading. The uncertainty bounds used in the tube design may therefore require periodic recalibration. Fleet data can be used to detect drift in disturbance characteristics and to update tube parameters in a principled manner. However, continuous data collection from vehicle fleets creates energy and storage demands, as well as privacy and data governance obligations. Sustainable deployment requires balancing the benefits of data-driven adaptation against the environmental and social costs of large-scale data infrastructure.

Hardware obsolescence presents another lifecycle challenge. Embedded control platforms may remain in service for many years, while algorithmic and solver developments progress rapidly. Event-triggered tube model predictive control can be implemented on modest hardware relative to more computationally intensive methods, but it still depends on numerical linear algebra routines, memory bandwidth, and real-time operating system support. Designing the controller so that it can be ported across hardware generations with minimal

revalidation effort is important for lifecycle sustainability. Likewise, the control architecture should avoid unnecessary coupling between the event-triggering policy and proprietary hardware features.

The broader sustainability implications of vehicle automation also intersect with motion control. Robust tracking that reduces unnecessary acceleration and braking can improve energy efficiency and reduce tire and brake wear. A controller that responds smoothly to disturbances may lower peak forces on suspension components, contributing to longer component life. Event-triggered operation can be designed to avoid rapid successive replanning events that lead to jerky motion or inefficient actuation. These physical effects are difficult to capture in abstract control theory but are significant in fleet-level environmental impact assessments.

9. Case Illustrations and Cross-Domain Comparisons

Consider a highway driving scenario in which a vehicle is exposed to periodic crosswinds as it passes bridge sections and open terrain. A time-triggered tube model predictive controller would replan at every sampling interval, even when the crosswind disturbance remains constant and the current nominal trajectory continues to provide safe guidance. In contrast, an event-triggered implementation can maintain the existing plan until the vehicle state approaches the edge of the robustness tube. The triggering rule may activate more frequently during high-curvature transitions or sudden gusts, but it remains relatively inactive during straight sections with steady disturbances. This behavior illustrates the core value of event triggering: computational effort is directed toward moments when replanning is truly needed.

A second scenario involves urban low-speed operation with frequent road irregularities and occluded localization landmarks. Here, sensor and localization uncertainty may be larger than dynamic uncertainty. The tube width must be expanded to account for intermittent positioning errors, and the triggering policy must be sensitive to the growth of state uncertainty. Data-driven methods that transfer models across urban environments can help estimate local disturbance characteristics before the vehicle enters a new road segment [6]. The event-triggered tube predictive control architecture can use such predictions to adjust tube width and triggering thresholds proactively, thereby avoiding sudden bursts of replanning when entering a previously unseen environment.

Cross-domain comparisons with networked industrial control illustrate the generality of the architectural principles. In process control, communication constraints and computational load often drive the adoption of event-triggered strategies. The same conceptual mechanisms appear in building energy management, robotic manipulation, and aerial vehicle control, where uncertain environments and limited onboard resources create similar tensions. The vehicle domain is distinctive because safety margins are small, actuator authority is constrained, and human life is directly affected by control failures. Nevertheless, the separation of nominal planning, disturbance rejection, and event-based execution can inform resilient control design in other sectors.

The required event-triggered tube predictive control formulation with input-to-state stability guarantees [18] can be compared with safety barrier approaches that impose state constraints through control inputs [17]. While barrier methods directly enforce forward invariance of safe sets, tube methods enforce safety through constraint tightening and bounded deviation from a nominal plan. The two approaches are not mutually exclusive and can be combined in a layered safety architecture. The comparison highlights a broader systems principle: no single

control mechanism should bear the entire burden of safety. Robustness is best achieved through complementary mechanisms that operate at different timescales and levels of abstraction.

10. Future Directions

Future research should develop adaptive tube geometries that respond to time-varying uncertainty estimates while preserving formal guarantees. Current tube methods often assume fixed disturbance bounds, but vehicle uncertainty changes with road type, weather, and mechanical condition. Learned disturbance models can provide context-dependent uncertainty predictions, but they must be integrated with rigorous safety envelopes. Safe learning methods offer a promising path, particularly when they allow the tube to shrink in well-characterized environments and expand cautiously in novel ones [16].

Event-triggering policies should also become more context-aware. A triggering decision should consider not only the current tracking error but also road curvature preview, traffic conditions, sensor health, and computational load. This will require a shift from static thresholds to more sophisticated triggering logic that can be validated using reachability analysis and data-driven scenario generation [10], [19]. At the same time, verification methods must evolve to address the aperiodic, data-dependent behavior of event-triggered systems. Statistical validation alone is unlikely to be sufficient for safety-critical certification, but exhaustive reachability analysis over all possible triggering sequences may be computationally prohibitive.

Finally, the governance and infrastructure dimensions require further interdisciplinary study. As vehicle control becomes more connected and data-driven, the boundaries between control, communication, and data infrastructure will blur. Cross-domain transfer methods [6] will increasingly influence control parameters, but institutional mechanisms for data quality assurance, liability assignment, and model drift monitoring remain underdeveloped. The future of robust vehicle motion tracking lies not only in better algorithms but also in the creation of governance frameworks and engineering practices that sustain robustness across the entire lifecycle.

11. Conclusion

Event-triggered tube model predictive control offers a principled systems architecture for robust vehicle motion tracking under uncertain dynamics and external disturbances. By combining tube-based uncertainty propagation with event-driven replanning, it reduces computational and communication load while preserving bounded tracking error and constraint satisfaction. The architecture separates nominal planning from disturbance rejection, enabling designers to reason explicitly about uncertainty, triggering thresholds, and real-time implementation constraints. This separation is valuable in safety-critical vehicle platforms where resources are limited and performance must be predictable.

However, the adoption of this architecture introduces structural, organizational, and governance challenges that extend beyond the control law. Computational scheduling becomes less regular, communication patterns become bursty, and certification must address aperiodic behavior. Data-driven adaptation can improve performance but also raises concerns about fairness, representativeness, and lifecycle maintenance. Sustainable deployment requires attention to embedded hardware capabilities, energy consumption, software portability, and fleet-level data governance. The event-triggered tube predictive control

framework is best understood as a socio-technical system in which formal robustness guarantees interact with infrastructure, policy, and operational practice.

A systems-level perspective reveals that the value of event-triggered tube model predictive control lies not in eliminating uncertainty but in making uncertainty an explicit, manageable component of the control architecture. This orientation supports robust, resource-aware, and safety-conscious vehicle motion tracking and provides a foundation for future developments in adaptive, data-informed, and certifiable autonomous vehicle control.

References

1. Mayne, D. Q., Rawlings, J. B., Rao, C. V., & Scokaert, P. O. M. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, 36(6), 789–814.
2. Rawlings, J. B., Mayne, D. Q., & Diehl, M. M. (2017). *Model predictive control: Theory, computation, and design* (2nd ed.). Nob Hill Publishing.
3. Langson, W., Chryssochoos, I., Raković, S. V., & Mayne, D. Q. (2004). Robust model predictive control using tubes. *Automatica*, 40(1), 125–133.
4. Tabuada, P. (2007). Event-triggered real-time scheduling of stabilizing control tasks. *IEEE Transactions on Automatic Control*, 52(9), 1680–1685.
5. Heemels, W. P. M. H., Johansson, K. H., & Tabuada, P. (2012). An introduction to event-triggered and self-triggered control. In *Proceedings of the 51st IEEE Conference on Decision and Control (CDC)*, 3270–3285.
6. Zhang, C., Zhang, H., Qiao, J., Yuan, D., & Zhang, M. (2019). Deep transfer learning for intelligent cellular traffic prediction based on cross-domain big data. *IEEE Journal on Selected Areas in Communications*, 37(6), 1389–1401.
7. Falcone, P., Borrelli, F., Asgari, J., Tseng, H. E., & Hrovat, D. (2007). Predictive active steering control for autonomous vehicle systems. *IEEE Transactions on Control Systems Technology*, 15(3), 566–580.
8. Gao, Y., Gray, A., Tseng, H. E., & Borrelli, F. (2014). A tube-based robust nonlinear predictive control approach to semiautonomous ground vehicles. *Vehicle System Dynamics*, 52(6), 802–823.
9. Paden, B., Čáp, M., Yong, S. Z., Yershov, D., & Frazzoli, E. (2016). A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Transactions on Intelligent Vehicles*, 1(1), 33–55.
10. Althoff, M., & Dolan, J. M. (2014). Online verification of automated road vehicles using reachability analysis. *IEEE Transactions on Robotics*, 30(4), 903–918.
11. Åström, K. J., & Wittenmark, B. (2013). *Computer-controlled systems: Theory and design* (3rd ed.). Dover Publications.
12. Maciejowski, J. M. (2002). *Predictive control with constraints*. Prentice Hall.
13. Mattingley, J., Wang, Y., & Boyd, S. (2011). Receding horizon control: Automatic generation of high-speed solvers. *IEEE Control Systems Magazine*, 31(3), 52–65.
14. Domahidi, A., Zgraggen, A. U., Zeilinger, M. N., Morari, M., & Jones, C. N. (2012). Efficient interior point methods for multistage problems arising in receding horizon control. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 668–674.

15. Chisci, L., Rossiter, J. A., & Zappa, G. (2001). Systems with persistent disturbances: Predictive control with restricted constraints. *Automatica*, 37(7), 1019–1028.
16. Hewing, L., Wabersich, K. P., Menner, M., & Zeilinger, M. N. (2020). Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3, 269–296.
17. Ames, A. D., Coogan, S., Egerstedt, M., Notomista, G., Sreenath, K., & Tabuada, P. (2019). Control barrier functions: Theory and applications. In 2019 18th European Control Conference (ECC), 3420–3431.
18. Cai, G., Huang, C., Feng, J., Zhou, Y., & Yin, G. (2026). Real-Time Event-Triggered Tube MPC With ISS Guarantees for Vehicle Motion Control Under Disturbances. *IEEE/ASME Transactions on Mechatronics*.
19. Kurzhanski, A. B., & Varaiya, P. (2000). Ellipsoidal techniques for reachability analysis: Part I: External approximations. In *Hybrid Systems: Computation and Control*, LNCS 1790, 202–214. Springer.
20. Sontag, E. D. (2008). Input to state stability: Basic concepts and results. In *Nonlinear and Optimal Control Theory*, Lecture Notes in Mathematics, 1932, 163–220. Springer.