Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Distributionally Robust Model Predictive Control with Disturbance Uncertainty Adaptiveness for Obstacle Avoidance of Autonomous Driving

Nan Xue
BA TechWorks, BMW Group
Lei Zhang
reizhg@tongji.edu.cn

Tongji University

Research Article

Keywords:

Posted Date: November 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9989197/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Distributionally Robust Model Predictive Control with Disturbance Uncertainty Adaptiveness for Obstacle Avoidance of Autonomous Driving

Nan Xue[1] and Lei Zhang[2*]

[1]BA TechWorks, BMW Group, Nanjing, China.
[2*]State Key Laboratory of High-speed Maglev Transportation Technology, College of Electronic and Information Engineering, Tongji University, Shanghai, China.

*Corresponding author(s). E-mail(s): reizhg@tongji.edu.cn;

Abstract

Traditionally, vehicle obstacle avoidance relies on nonlinear model predictive control (NMPC), which suffers from two major drawbacks. First, NMPC incurs higher computational costs compared to its linear counterpart. Second, it may fail in complex scenarios due to simplified modeling of time-varying vehicle dynamics and the presence of noisy disturbances. In this paper, we propose a distributionally robust model predictive control (MPC) framework for autonomous vehicle obstacle avoidance under nonlinear dynamics and unknown disturbance distributions. The Koopman operator is employed to transform the nonlinear vehicle dynamics into a linear representation in a high-dimensional latent space. Distributionally robust optimization based on a Wasserstein ambiguity set, constructed from historical data, is introduced to guarantee robustness against distributional shifts. The robust probabilistic obstacle avoidance constraint is further converted into tractable linear constraints. Averaged simulations on a racing circuit demonstrate that the proposed method preserves nominal tracking performance while improving feasibility and tail-error robustness under disturbance shifts, compared with non-robust Koopman MPC.

Keywords: Vehicle Dynamics, Distributionally Robust Model Predictive Control, Koopman Operator, Uncertainty Adaptiveness

1 INTRODUCTION

1.1 Research Background

Decision-making and control constitute the core of autonomous driving technology, directly determining vehicle performance, safety, and comfort. However, the increasing complexity of driving scenarios—ranging from urban congestion to high-speed racing—poses significant challenges to conventional control approaches. In extreme dynamic scenarios such as real-time racing control, there are exceptionally high demands on the system's decision-making and control capabilities. The safe deployment in complex real-world environments faces two key challenges: the inherent nonlinearity of vehicle dynamics and the unknown or mismatched distribution of external disturbances. In general driving scenarios, vehicles frequently operate under strongly nonlinear dynamics during maneuvers such as high-speed racing, emergency obstacle avoidance, or abrupt lane changes. Simultaneously, factors such as variations in road friction coefficients, crosswind disturbances, and uncertain behaviors of surrounding traffic participants make it difficult to accurately model the statistical properties of disturbances.

To manage the computational complexity of NMPC, the nonlinear vehicle dynamics are typically linearized. However, traditional linearization techniques are valid only near equilibrium points and fail to capture global vehicle behavior during safety-critical maneuvers such as real-time obstacle avoidance. Conventional stochastic control relies on known disturbance distribution assumptions, which can lead to catastrophic outcomes when the distribution is mismatched. While deterministic robust control can guarantee safety in the worst-case scenario, it often proves too conservative to meet the extreme performance demands of driving.

In recent years, Koopman operator theory has offered a new perspective for addressing nonlinear control problems. By lifting the state to a high-dimensional latent space, it enables the global linearization of nonlinear systems, thereby allowing linear control methods to be applied while preserving dynamic integrity. Concurrently, distributionally robust model predictive control (DR-MPC) characterizes disturbance uncertainty by constructing an ambiguity set that contains possible true distributions. It considers the worst-case distribution within the optimization objective, effectively balancing performance and robustness. Nevertheless, the systematic integration of these two methods for nonlinear vehicle obstacle avoidance control—particularly regarding the tractability of probabilistic safety constraints in real-time racing scenarios—still requires further investigation.

Existing works have explored Koopman-based MPC [6, 7, 15] and distributionally robust control [9–11] separately; the integration of these two paradigms—especially with probability constraints imposed directly on the lifted disturbances—remains unexplored. This paper fills this gap by proposing a unified framework that combines the linearization power of the deep Koopman operator with the robustness guarantees of Wasserstein distributionally robust optimization (DRO) to enhance safety performance in vehicle obstacle avoidance.

Specifically, we propose a distributionally robust MPC strategy with uncertainty adaptiveness. The framework first employs the Koopman operator to achieve global

linearization of the nonlinear vehicle dynamics. It then uses the Wasserstein distance to construct an ambiguity set for the disturbance distribution, capturing distributional uncertainty within a minimax optimization framework. Finally, probabilistic obstacle avoidance constraints are equivalently transformed into tractable chance constraints, ensuring high-probability satisfaction of safety constraints under uncertainty. Extensive aggregated simulations on nonlinear vehicle models demonstrate that the proposed method improves severe-shift feasibility and tail-error robustness compared to non-robust Koopman MPC, while retaining comparable nominal performance.

1.2 Related Works

Our work lies at the intersection of Koopman operator theory, distributionally robust optimization (DRO), and model predictive control (MPC), aiming to achieve obstacle avoidance for autonomous vehicles. This section reviews the most relevant advancements in each field and positions our contribution.

The Koopman operator provides a linear representation of nonlinear dynamics by lifting the original states to a space of observable functions [1]. Its data-driven approximation via Extended Dynamic Mode Decomposition (EDMD) was introduced by Williams et al. [2], while Lusch et al. [3] and Shi and Meng [4] later proposed deep learning-based lifting functions (Deep Koopman) for improved expressiveness. In the context of autonomous vehicles, Xiao et al. [8] applied Deep Koopman to modeling and control, and Wu et al. [5] used it for freeway ramp metering. Zhang et al. [7] developed a robust tube-based Koopman MPC for nonlinear systems. Zheng et al. [6] proposed a Koopman-based hybrid modeling and zonotopic tube robust MPC for automated vehicle motion control, where the Koopman modeling error is explicitly treated as a disturbance. More recently, Chen and Lv [22] introduced an extended state observer (ESO) to compensate for Koopman modeling residuals online, while Abtahi et al. [21] proposed a deep bilinear Koopman model in the Frenet frame and validated it through hardware-in-the-loop experiments on a dSPACE platform. Related racing-oriented trajectory estimation with composite Bezier curves was further reported in Weiss and Behl [26]. These works prove the effectiveness of Koopman-based linearization but typically assume perfect knowledge of the system model or neglect distributional uncertainty in disturbances.

Control based on Distributionally Robust Optimization (DRO) has emerged as a powerful tool to handle unknown disturbance distributions. Esfahani and Kuhn [9] laid the foundation for data-driven DRO using the Wasserstein metric, providing finite-sample guarantees and tractable reformulations. Li et al. [10] incorporated DRO into stochastic MPC, while Li et al. [16] extended it to output-feedback settings. For motion control, Hakobyan and Yang [11] proposed a Wasserstein DRO framework with distributionally robust chance constraints for vehicle collision avoidance. Chen et al. [12] applied DRO to motion planning for autonomous driving, and related MPC-based collision-avoidance analysis in distributed robot systems was studied in Vangasse et al. [24]. Most recently, Fang et al. [23] proposed a distributionally robust chance-constrained linear parameter-varying MPC (DRCC-LPVMPC) that treats model mismatch as additive uncertainty with unknown distribution, constructing Wasserstein ambiguity sets from finite samples and validates on real vehicles.

Despite these advances, most DRO-based controllers apply robustness to state or output constraints, rather than to the underlying disturbance itself.

The integration of the Koopman operator with MPC has been explored by Zhang et al. [15] and Zhang et al. [7]. From a different modeling perspective, Nezami et al. [27] developed an LPVMPC obstacle-avoidance framework that introduces a scheduling trust region to preserve feasibility under scheduling mismatch. In the autonomous racing domain, Xue et al. [29] combined learning-based error dynamics with MPC for high-speed trajectory tracking, Numerow et al. [30] proposed an inherently robust MPC formulation with anytime feasible SQP for real-time racing, and Zhu and Borrelli [28] formulated multi-agent racing interactions as open-loop generalized Nash equilibria. Practical comparisons between data-enabled predictive control and model-based predictive control for robotic systems have also been reported in Chen et al. [25]. However, existing Koopman-MPC frameworks typically characterize uncertainties using only first and second moments, an assumption that fails to capture the full complexity of real-world driving disturbances. Moreover, none of these works impose robustness directly on the lifted disturbance within the Koopman latent space, nor do they leverage Wasserstein chance-constraint dual reformulations to bound disturbance energy.

In contrast, our proposed Koopman-based DR-MPC method achieves global linearization in a high-dimensional latent space and adopts Wasserstein ambiguity sets to model disturbance uncertainty. This avoids both the approximation errors of local linearization methods and the restrictive moment assumptions of prior Koopman-DRMPC frameworks, while maintaining computational tractability for safety-critical obstacle avoidance.

1.3 Contributions

The contributions of our paper are presented as follows:

i) A Koopman-based linearization framework is designed for real-time obstacle avoidance. The Koopman operator is employed to lift inherently nonlinear vehicle dynamics into a high-dimensional latent space, achieving global linearization that breaks the local validity limitation of traditional linearization methods. It is then implemented via Extended Dynamic Mode Decomposition (EDMD) combined with deep neural networks (Deep Koopman), enabling direct application of linear control theories and improving modeling accuracy for fast obstacle avoidance.

ii) A safety-oriented distributionally robust MPC (DR-MPC) strategy is proposed to address unknown or mismatched disturbance distributions. A Wasserstein ambiguity set is defined to characterize disturbance uncertainty, and the control problem is cast as a min-max optimization to minimize the worst-case expected cost over plausible distributions to avoid the over-conservatism of deterministic robust MPC and strong distribution assumptions of stochastic MPC.

iii) An equivalent reformulation is presented for robust probabilistic obstacle avoidance constraints. Leveraging Wasserstein dual reformulations of chance constraints, the intractable probabilistic constraint is converted into a set of linear constraints, which are solvable via standard convex solvers to meet real-time control requirements.

iv) Extensive repeated simulations are conducted to validate the proposed framework under distribution shifts. Results show improved severe-shift feasibility and better tail-error robustness relative to non-robust Koopman MPC, while preserving comparable nominal performance.

Notation

The following notation is used throughout the paper.

- $s_t \in \mathbb{R}^5$: physical state vector $s_t = [p_{x,t},\ p_{y,t},\ \psi_t,\ v_t,\ \omega_t]^\top$, where (p_x, p_y) are global coordinates, ψ is yaw angle, v is longitudinal velocity, ω is yaw rate.
- $u_t \in \mathbb{R}^2$: control input $u_t = [a_t,\ \delta_t]^\top$, where a is acceleration and δ is steering angle.
- $w_t \in \mathbb{R}^5$: external physical disturbance (road irregularities, wind, etc.).
- $z_t \in \mathbb{R}^{n_z}$: lifted state (Koopman latent space).
- ψ^e, ψ^d: encoder and decoder network.
- $A \in \mathbb{R}^{n_z \times n_z}$, $B \in \mathbb{R}^{n_z \times 2}$: Koopman linear dynamics matrices.
- $\hat{w}_t(s_t, u_t, w_t) \in \mathbb{R}^{n_z}$: combined disturbance (model mismatch + external effects), depends on state, control, and external noise.
- $D \in \mathbb{R}^{2\times 5}$, $E \in \mathbb{R}^{1\times 5}$, $F \in \mathbb{R}^{1\times 5}$: linear projection matrices from latent state to position, velocity, yaw rate.
- $G_{\text{road},t} \in \mathbb{R}^{2\times 2}$, $h_{\text{road},t} \in \mathbb{R}^2$: stagewise road-boundary half-space coefficients obtained from tangent lines to the left and right road edges.
- $H_{\text{obs},j} \in \mathbb{R}^{1\times 2}$, $h_{\text{obs},j} \in \mathbb{R}$: linear inequality coefficients used to enforce a rectangular obstacle-avoidance half-space for obstacle j.
- $Q = \text{diag}(q_x, q_y, q_\psi, q_v, q_\omega)$, $R = \text{diag}(r_a, r_\delta)$: weight matrices for cost function.
- T: prediction horizon; Δt: sampling time.
- $\eta > 0$: safety margin for obstacle avoidance (small distance).
- $\epsilon \in (0, 1)$: risk level for chance constraints (e.g., 0.05).
- $\theta \geq 0$: radius of the Wasserstein ambiguity set.
- $\{z_t, u_t\}_{t=0}^{T-1}$: lifted state and control sequences over the prediction horizon.

2 VEHICLE MODEL AND OBSTACLE AVOIDANCE

2.1 Problem Formulation

2.1.1 Dynamic Model

To describe the vehicle motion at moderate to high speeds, we consider the kinematic bicycle model, which captures the essential geometric relationship between velocity, steering angle, and vehicle orientation without relying on tire force equations.

$$s_{t+1} = f(s_t, u_t, w_t), \tag{2.1}$$

This equation describes the evolution of the vehicle from time step t to $t+1$ under the nonlinear dynamics $f(\cdot)$, which is defined by the kinematic bicycle model.

The variables are defined as:

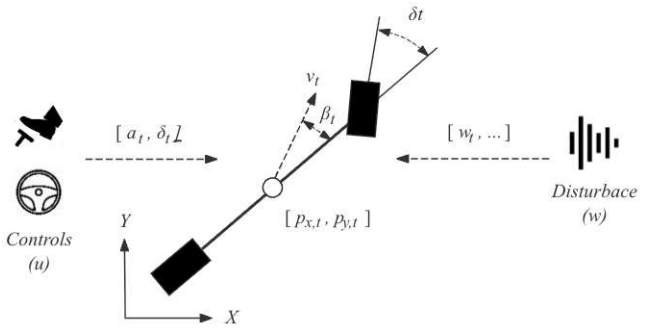


Fig. 2.1 Vehicle dynamics (kinematic bicycle model).

i) $s_t \in \mathbb{R}^5$ denotes the state vector, defined as:

$$s_t = \begin{bmatrix} p_{x,t} & p_{y,t} & \psi_t & v_t & \omega_t \end{bmatrix}^\top, \tag{2.2}$$

where $t \in \{0, 1, \ldots, T\}$ are discrete time steps, p_x and p_y are the global position coordinates, ψ is the yaw angle, v is the longitudinal velocity, and ω is the yaw rate.

ii) $u_t \in \mathbb{R}^2$ denotes the control input, defined as:

$$u_t = \begin{bmatrix} a_t & \delta_t \end{bmatrix}^\top, \tag{2.3}$$

where a is the acceleration (positive for throttle, negative for braking) and δ is the steering angle.

iii) $w_t \in \mathbb{R}^5$ denotes the disturbance vector, defined as:

$$w_t = \begin{bmatrix} w_{x,t} & w_{y,t} & w_{\psi,t} & w_{v,t} & w_{\omega,t} \end{bmatrix}^\top, \tag{2.4}$$

capturing the effect of exogenous uncertainties affecting the vehicle dynamics, including road irregularities, wind gusts, tire-road friction variations, and model mismatches between the nominal dynamics and the real system.

2.1.2 Optimization Objective

The objective is to track a reference trajectory while minimizing control effort:

$$\min_{\{u_t\}_{t=0}^{T-1}} \sum_{t=0}^{T-1} \left(\|s_t - s_t^{\text{ref}}\|_Q^2 + \|\Delta u_t\|_R^2 \right), \tag{2.5}$$

where $\Delta u_t = u_t - u_{t-1}$.

The weighting matrices are defined as:

$$\begin{aligned} Q &= \mathrm{diag}(q_x,\ q_y,\ q_\psi,\ q_v,\ q_\omega) \in \mathbb{R}^{5\times 5}, \\ R &= \mathrm{diag}(r_a,\ r_\delta) \in \mathbb{R}^{2\times 2}, \end{aligned} \tag{2.6}$$

with positive scalars $q_x, q_y, q_\psi, q_v, q_\omega$ penalizing the tracking errors in position, yaw angle, longitudinal velocity, and yaw rate, respectively, and r_a, r_δ penalizing changes in acceleration and steering angle. The squared Mahalanobis distance is defined as $\|s\|_Q^2 = s^\top Q s$.

2.1.3 State Constraints

The vehicle must satisfy the following physical and safety limits at each time step. We retain a rectangular obstacle description for surrounding vehicles and represent the local drivable corridor by tangent half-spaces constructed from the road edges.

i) Obstacle avoidance constraint (rectangular): Each surrounding obstacle j is represented by an axis-aligned rectangle $[x_j^{\min}, x_j^{\max}] \times [y_j^{\min}, y_j^{\max}]$. To preserve convexity and real-time tractability, obstacle avoidance is enforced through a preselected separating half-space

$$H_{\mathrm{obs},j} \begin{bmatrix} p_{x,t} \\ p_{y,t} \end{bmatrix} \le h_{\mathrm{obs},j}, \quad \forall j, \forall t \in \mathbb{Z}_{0:T}. \tag{2.7}$$

The pair $(H_{\mathrm{obs},j}, h_{\mathrm{obs},j})$ is selected according to the relative position between the ego vehicle and the obstacle before solving the MPC problem, so that the vehicle is constrained to pass the obstacle on the chosen safe side with margin $\eta > 0$.

ii) Road boundary constraint: At each prediction step, the local drivable corridor is approximated by two tangent lines constructed from the left and right road boundaries around the predicted reference point. Each tangent line is written in the standard form $ax + by = c$, where (a, b) is the inward normal vector of the corresponding road edge and c is determined by the boundary point through which that tangent line passes. Therefore, $(a_{1,t}, b_{1,t}, c_{1,t})$ and $(a_{2,t}, b_{2,t}, c_{2,t})$ denote the line coefficients associated with the left and right road boundaries at prediction step t, respectively. To write both boundary conditions in the same "$\le$" form required by the optimizer, we define

$$\begin{aligned} G_{\mathrm{road},t} &= \begin{bmatrix} a_{1,t} & b_{1,t} \\ -a_{2,t} & -b_{2,t} \end{bmatrix}, \\ h_{\mathrm{road},t} &= \begin{bmatrix} c_{1,t} \\ -c_{2,t} \end{bmatrix}, \end{aligned} \tag{2.8}$$

so that the left-boundary constraint $a_{1,t}p_{x,t} + b_{1,t}p_{y,t} \le c_{1,t}$ and the right-boundary constraint $a_{2,t}p_{x,t} + b_{2,t}p_{y,t} \ge c_{2,t}$ are represented jointly. Then the vehicle position

must satisfy

$$G_{\text{road},t} \begin{bmatrix} p_{x,t} \\ p_{y,t} \end{bmatrix} \leq h_{\text{road},t}, \quad \forall t \in \mathbb{Z}_{0:T}. \tag{2.9}$$

This stagewise representation keeps the vehicle inside the local road corridor while adapting the admissible half-spaces along the curved track.

iii) Longitudinal velocity limits: The vehicle's forward speed is bounded by physical and legal limits:

$$v_{\min} \leq v_t \leq v_{\max}, \quad \forall t. \tag{2.10}$$

iv) Acceleration limits: The longitudinal acceleration is bounded by actuator capabilities:

$$a_{\min} \leq a_t \leq a_{\max}, \quad \forall t. \tag{2.11}$$

v) Yaw rate limits: The yaw rate is constrained to prevent excessive rotation:

$$\omega_{\min} \leq \omega_t \leq \omega_{\max}, \quad \forall t. \tag{2.12}$$

2.1.4 Control Constraints

The control input set $\mathcal{U} \subset \mathbb{R}^2$ reflects actuator limits:

$$\mathcal{U} = \left\{ u = [a, \delta]^\top \in \mathbb{R}^2 \;\middle|\; a_{\min} \leq a \leq a_{\max},\ |\delta| \leq \delta_{\max} \right\}. \tag{2.13}$$

The steering rate is additionally constrained by actuator dynamics:

$$|\delta_{t+1} - \delta_t| \leq \Delta\delta_{\max}\Delta t, \quad \forall t \in \mathbb{Z}_{0:T-2}. \tag{2.14}$$

2.1.5 Nominal MPC Problem

The nominal MPC problem is formulated as:

$$\begin{aligned}
\min_{\{u_t\}_{t=0}^{T-1}} \quad & \sum_{t=0}^{T-1} \left(\|s_t - s_t^{\text{ref}}\|_Q^2 + \|\Delta u_t\|_R^2 \right) \\
\text{s.t.} \quad & s_{t+1} = f(s_t, u_t) + w_t, \quad \forall t \in \mathbb{Z}_{0:T-1}, && \text{(2.15a)} \\
& s_0 = s_{\text{current}}, && \text{(2.15b)} \\
& v_{\min} \leq v_t \leq v_{\max}, \quad \forall t \in \mathbb{Z}_{0:T}, && \text{(2.15c)} \\
& \omega_{\min} \leq \omega_t \leq \omega_{\max}, \quad \forall t \in \mathbb{Z}_{0:T}, && \text{(2.15d)} \\
& a_{\min} \leq a_t \leq a_{\max}, \quad \forall t \in \mathbb{Z}_{0:T-1}, && \text{(2.15e)} \\
& |\delta_t| \leq \delta_{\max}, \quad \forall t \in \mathbb{Z}_{0:T-1}, && \text{(2.15f)} \\
& |\delta_{t+1} - \delta_t| \leq \Delta\delta_{\max}\Delta t, \quad \forall t \in \mathbb{Z}_{0:T-2}, && \text{(2.15g)} \\
& G_{\text{road},t} \begin{bmatrix} p_{x,t} \\ p_{y,t} \end{bmatrix} \leq h_{\text{road},t}, \quad \forall t \in \mathbb{Z}_{0:T}, && \text{(2.15h)} \\
& H_{\text{obs},j} \begin{bmatrix} p_{x,t} \\ p_{y,t} \end{bmatrix} \leq h_{\text{obs},j}, \quad \forall j, \forall t \in \mathbb{Z}_{0:T}. && \text{(2.15i)}
\end{aligned}$$

Here, the obstacle constraint in (2.15i) keeps the rectangular avoidance rule in a linear half-space form, while the road constraint in (2.15h) is represented by time-varying tangent half-spaces. This formulation serves as the baseline for the Koopman lifting and the tractable robust reformulations developed in the subsequent sections.

3 DATA-DRIVEN DEEP KOOPMAN MODEL

In this section, we will introduce the Koopman operator, the extended dynamic mode decomposition (EDMD) method, and a data-driven Deep Koopman model for autonomous vehicle obstacle avoidance.

3.1 The Koopman Operator

The Koopman operator was originally developed to describe the inherent properties of an uncontrolled nonlinear dynamical system through a linear evolutionary framework. With a slight modification, it can be adapted to controlled dynamical systems. In this paper, we consider a nonlinear dynamical system modeling autonomous vehicle dynamics. For clarity, we omit the noise term from system (2.1), yielding the following deterministic form:

$$s_{t+1} = f(s_t, u_t), \tag{3.1}$$

Let $\mathbf{u}_t = (u_t, u_{t+1}, u_{t+2}, \ldots) \in \bar{\mathcal{U}}$, and define $\mathcal{G}$ as the left shift operator such that $\mathcal{G}\mathbf{u}_t = \mathbf{u}_{t+1}$. Here, $\bar{\mathcal{U}}$ denotes the space of all sequences $(u_0, u_1, u_2, \ldots)$ with $u_t \in \mathcal{U}$. We then define the Koopman operator for the dynamics on the extended state $[s, \mathbf{u}]$ as follows:

$$\mathcal{K}g(s_t, \mathbf{u}_t) = g(s_{t+1}, \mathbf{u}_{t+1}) = g(f(s_t, u_t), \mathcal{G}\mathbf{u}_t), \tag{3.2}$$

where $\mathcal{K} : \mathcal{H} \to \mathcal{H}$ is the Koopman operator, $g : \mathcal{X} \times \bar{\mathcal{U}} \to \mathbb{R}^N$ is an observable function belonging to $\mathcal{H}$, and $\mathcal{H}$ is a Banach space. The Koopman operator $\mathcal{K}$ acts as a composition of g with f. Notably, $\mathcal{K}$ is infinite-dimensional and linear, even when the dynamics $f(\cdot)$ is nonlinear. Due to this linearity, $\mathcal{K}$ admits a standard eigendecomposition:

$$\mathcal{K}\phi_j(s_t, \mathbf{u}_t) = \lambda_j \phi_j(s_t, \mathbf{u}_t), \quad j = 1, 2, \ldots, \tag{3.3}$$

where $\phi_j(s_t, \mathbf{u}_t)$ and λ_j are the eigenfunction and associated eigenvalue of $\mathcal{K}$, respectively. A vector-valued observable g can be expressed in terms of these eigenfunctions as:

$$g(s_t, \mathbf{u}_t) = (\mathcal{K}^t g)(s_0, \mathbf{u}_0) = \sum_{j=1}^{\infty} \lambda_j^t \phi_j(s_0, \mathbf{u}_0)\nu_j, \tag{3.4}$$

where $\mathcal{K}^t$ denotes the t-fold application of the Koopman operator, ν_j is the Koopman mode associated with eigenvalue λ_j, and s_0 is the initial state.

3.2 Extended Dynamic Mode Decomposition

The extended dynamic mode decomposition (EDMD) is a method for obtaining a finite-dimensional approximation of the Koopman operator. EDMD employs a set of basis functions—such as radial basis functions (RBFs) with varying kernel centers and

widths—to represent the observable functions, and uses the least squares method to compute a finite-dimensional approximation $\mathbf{K} \in \mathbb{R}^{N\times N}$ of $\mathcal{K}$. Note that it is generally infeasible to directly apply the least squares method to compute an approximation of the Koopman operator in finite time, because the extended state $[s, \mathbf{u}]$ is typically infinite-dimensional.

To obtain the finite-dimensional approximation $\mathbf{K}$, we choose the lifting functions g_i for $i = 1, 2, \ldots, N$ in a specific form:

$$g_i(s_t, \mathbf{u}_t) = \psi_i(s_t) + h_i(\mathbf{u}_t), \tag{3.5}$$

where $\psi_i : \mathcal{X} \to \mathbb{R}$ is generally nonlinear, while $h_i : \bar{\mathcal{U}} \to \mathbb{R}$ is linear. Without loss of generality, we assume the vector-valued observable function $g = [g_1, g_2, \ldots, g_N]^\top$ takes the form:

$$g(s_t, \mathbf{u}_t) = [\psi^\top(s_t), u_t^\top]^\top, \tag{3.6}$$

where $\psi(s_t) = [\psi_1(s_t), \psi_2(s_t), \ldots, \psi_l(s_t)]^\top$ with $l = N - m$, and $u_t \in \mathbb{R}^m$ is the first component of the sequence $\mathbf{u}_t$. To avoid notational abuse, we consistently use u_t to denote this first component throughout the paper. From (3.2) and (3.6), we obtain:

$$\mathcal{K}g(s_t, \mathbf{u}_t) = \mathbf{K}[\psi^\top(s_t), u_t^\top]^\top + r(s_t, u_t), \tag{3.7}$$

where $r(s_t, u_t)$ is a residual term representing the discrepancy between the N-dimensional approximation of the observable space and the true lifted space of the Koopman operator. We determine $\mathbf{K}$ by minimizing the cost function $\|r(s_t, u_t)\|_2^2$. Since the future values of the control sequence are not needed for prediction, we can discard the last m components of $[\psi^\top(s_{t+1}), u_{t+1}^\top]^\top$. Defining $\Psi(s_t) = [\psi^\top(s_t), u_t^\top]^\top$, we rewrite the observable dynamics governed by the Koopman operator as:

$$\begin{aligned} \psi(s_{t+1}) &= \mathbf{A}\Psi(s_t), \\ \tilde{s}_t &= C\psi(s_t), \end{aligned} \tag{3.8}$$

where $\mathbf{A} = [A, B] \in \mathbb{R}^{l\times N}$ consists of $A \in \mathbb{R}^{l\times l}$ and $B \in \mathbb{R}^{l\times m}$. Here, $\mathbf{A}$ corresponds to the first l rows of $\mathbf{K}$, $C \in \mathbb{R}^{n\times l}$ maps the lifted space back to the original state space, and $\tilde{s}_t \in \mathbb{R}^n$ is the reconstructed state approximating s_t. Consequently, the matrices A and B are obtained by solving the optimization problem:

$$\min_{A,B} \|r\|_2^2 = \min_{A,B} \sum_{t=0}^{T} \|\psi(s_{t+1}) - A\psi(s_t) - Bu_t\|_2^2. \tag{3.9}$$

3.3 The Deep Koopman Method

Compared with the classic EDMD, the Deep Koopman model employs a deep neural network to generate the observable subspace of the Koopman operator. The proposed Deep Koopman model adopts an encoder-decoder architecture, where the encoder ψ^e maps the system inputs to a lifted space, and the decoder ψ^d reconstructs the original

system states from the lifted features. The dynamical system is represented as:

$$\psi^e(s_{t+1}, w^e) = \widetilde{\mathbf{A}}\Phi(s_t, w^e, u_t), \tag{3.10}$$

$$\tilde{s}_t = \psi^d(\psi^e(s_t, w^e), w^d), \tag{3.11}$$

where $\widetilde{\mathbf{A}} = [\widetilde{A}, \widetilde{B}] \in \mathbb{R}^{l\times N}$, $\Phi(s_t, w^e, u_t) = [\psi^e(s_t, w^e)^\top, u_t^\top]^\top \in \mathbb{R}^N$, ψ^e is the encoder with parameters w^e, and ψ^d is the decoder with parameters w^d.

For the autonomous vehicle control problem, the state s exhibits temporal dependencies. Therefore, a long short-term memory (LSTM) network is used to process sequential data within the encoder, i.e., $\psi^e(s, w^e)$, which lifts the original state s to a higher-dimensional vector of observables.

Let K denote the sequence length of the state inputs to the encoder. The vector-valued observables are then given by the last hidden state of the LSTM:

$$\psi^e(s, w^e) = [s^\top, h_K^\top]^\top. \tag{3.12}$$

The decoder ψ^d is constructed using a feed-forward neural network consisting of K' fully connected layers, which recover the original state from the higher-dimensional observables. The output of any hidden layer $k \in \{1, 2, \ldots, K'\}$ is expressed as:

$$h_k^d = \sigma_k^d(W_k^d h_{k-1}^d + b_k^d). \tag{3.13}$$

The predicted state is obtained from the output of the decoder's final layer:

$$\tilde{s}_t = \psi^d(\psi^e(s_t, w^e), w^d) = \sigma(W_{K'}^d h_{K'-1}^d + b_{K'}^d). \tag{3.14}$$

Then, the parameters W and b are learned by optimizing a fixed loss function during training. To accurately approximate autonomous vehicle dynamics over a longer time horizon, a multi-step prediction error is considered instead of a single-step error. To formulate the multi-step loss function, we first present the state prediction over an M-step interval:

$$s_{t+M} = \psi^d(\widetilde{A}^M\Phi(s_t, w^e, u_t), w^d) + r(s, M), \tag{3.15}$$

where $\widetilde{A}^M\Phi(s_t, w^e, u_t)$ represents the M-step ahead lifted state starting from s_t, given by:

$$\begin{aligned}\widetilde{A}^M\Phi(s_t, w^e, u_t) &= \psi^e(s_{t+M}, w^e) \\ &= \widetilde{A}\psi^e(s_{t+M-1}, w^e) + \widetilde{B}u_{t+M-1} \\ &= \widetilde{A}^M\psi^e(s_t, w^e) + \sum_{s=1}^{M}\widetilde{A}^{s-1}\widetilde{B}u_{t+M-s}.\end{aligned} \tag{3.16}$$

The main components of the loss function over M time steps are defined as:

$$L_1 = \frac{1}{M}\sum_{s=1}^{M} \|s_{t+s} - \psi^d(\widetilde{A}^s \Phi(s_t, w^e, u_t), w^d)\|_2^2, \tag{3.17}$$

$$L_2 = \frac{1}{M}\sum_{s=1}^{M} \|\psi^e(s_{t+s}, w^e) - \widetilde{A}^s \Phi(s_t, w^e, u_t)\|_2^2, \tag{3.18}$$

$$L_3 = \frac{1}{M}\sum_{s=1}^{M} \|s_s - \psi^d(\psi^e(s_s, w^e), w^d)\|_2^2, \tag{3.19}$$

where L_1 denotes the prediction error, L_2 the prediction error in the lifted observable space, and L_3 the reconstruction error. Using these definitions, the overall optimization problem is formulated as:

$$\min_{w^e, w^d, \widetilde{A}, \widetilde{B}} L = \alpha_1 L_1 + \alpha_2 L_2 + \alpha_3 L_3 + \alpha_4 \|w^e\|_2^2 + \alpha_5 \|w^d\|_2^2, \tag{3.20}$$

where α_i, $i = 1, 2, \ldots, 5$, are weight parameters, and the last two terms serve as regularization terms to prevent overfitting.

After training the Deep Koopman model, the approximated dynamics for the autonomous vehicle are given by:

$$\psi^e(s_{t+1}, w^e) = \widetilde{A}\psi^e(s_t, w^e) + \widetilde{B}u_t, \tag{3.21}$$

$$\tilde{s}_t = \psi^d(\psi^e(s_t, w^e), w^d). \tag{3.22}$$

3.4 Distributionally Robust Obstacle Avoidance

To ensure safe obstacle avoidance under uncertain disturbances, we impose a distributionally robust chance constraint on the safety margin over the entire prediction horizon. For this purpose, after training the Deep Koopman model, we obtain the encoder ψ^e, the decoder ψ^d, and the linear dynamics matrices $\widetilde{A}$ and $\widetilde{B}$. The lifted state evolves according to:

$$z_{t+1} = \widetilde{A}z_t + \widetilde{B}u_t + \hat{w}(s_t, u_t, w_t), \quad z_t = \psi^e(s_t), \tag{3.23}$$

and the physical state is recovered via $s_t = \psi^d(z_t)$, where $\hat{w}(\cdot)$ captures the combined effect of model mismatch and external disturbances. In the Koopman lifting design, the first five components of the latent state z_t correspond directly to the original physical state:

$$z_t = [p_{x,t}, p_{y,t}, \psi_t, v_t, \omega_t, \ldots]^\top. \tag{3.24}$$

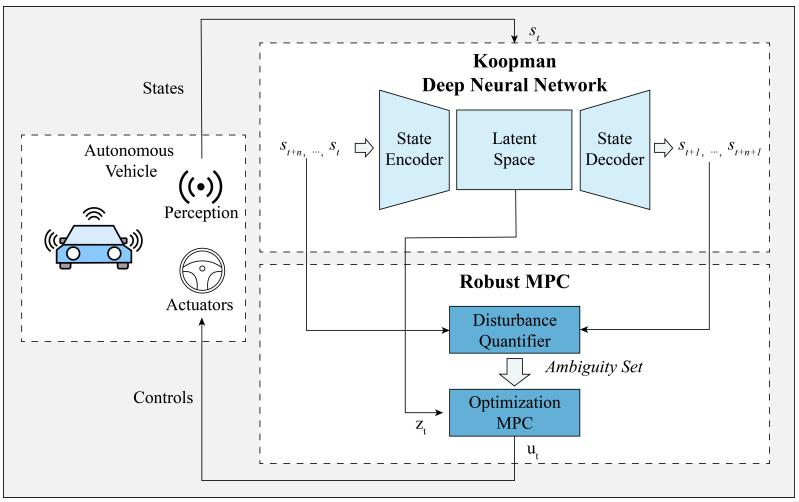


Fig. 3.1 Overall architecture of the proposed Koopman-based robust MPC framework. The measured vehicle states are lifted to a latent space through a deep Koopman encoder-decoder model. The latent dynamics and disturbance quantification are then integrated into a robust MPC optimization with a Wasserstein ambiguity set, and the resulting control input is fed back to the autonomous vehicle.

We treat the entire combined disturbance sequence over the prediction horizon as a single random vector. Define:

$$\hat{\mathbf{w}} = \begin{bmatrix} \hat{w}_0 \\ \hat{w}_1 \\ \vdots \\ \hat{w}_{T-1} \end{bmatrix} \in \mathbb{R}^{n_z T}, \tag{3.25}$$

where $\hat{w}_t = \hat{w}_t(s_t, u_t, w_t)$ is the combined disturbance at time t. Given N independent historical trajectories $\{\hat{\mathbf{w}}^{(i)}\}_{i=1}^{N}$, the empirical distribution is $\hat{\mathbb{P}}_N = \frac{1}{N}\sum_{i=1}^{N} \delta_{\hat{\mathbf{w}}^{(i)}}$. The Wasserstein ambiguity set of radius $\theta \geq 0$ is:

$$\mathcal{D} = \left\{ \mathbb{P} \in \mathcal{P}(\mathbb{R}^{n_z T}) \;\middle|\; W_1(\mathbb{P}, \hat{\mathbb{P}}_N) \leq \theta \right\}, \tag{3.26}$$

with W_1 denoting the first-order Wasserstein distance.

Define the linear projection matrices as follows:

$$D = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \end{bmatrix}, \quad E = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & \cdots & 0 \end{bmatrix}, \quad F = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & \cdots & 0 \end{bmatrix}. \tag{3.27}$$

Thus, the position, velocity, and yaw rate can be extracted as:

$$Dz_t = \begin{bmatrix} p_{x,t} \\ p_{y,t} \end{bmatrix}, \quad Ez_t = v_t, \quad Fz_t = \omega_t. \tag{3.28}$$

Let $l_{j,t}(s_t, \hat{w}_t)$ denote the safety margin for obstacle j at time t:

$$l_{j,t}(s_t, \hat{w}_t) = H_{\text{obs},j} D z_t - h_{\text{obs},j} \leq 0. \tag{3.29}$$

Then, distributionally robust obstacle avoidance is enforced via the following constraints:

$$\inf_{\mathbb{P}\in\mathcal{D}} \Pr_{\mathbb{P}}\big(l_{j,t}(s_t, \hat{w}_t) \leq 0\big) \geq 1 - \epsilon, \quad \forall j,\ \forall t = 0, \ldots, T-1. \tag{3.30}$$

Here $\epsilon \in (0, 1)$ is a small risk parameter (e.g., 0.05). Because the ambiguity set $\mathcal{D}$ contains distributions of the entire trajectory $\hat{\mathbf{w}}$, the marginal distribution of $\hat{w}_t$ is not fixed. Constraint (3.30) guarantees that for each time step, even under the worst-case joint distribution within $\mathcal{D}$, the probability of violating safety at that step is at most ϵ.

Under this setting, the stage cost is defined as:

$$c(z_t, u_t) = \|D z_t - y_t^{\text{ref}}\|_Q^2 + \|\Delta u_t\|_R^2, \tag{3.31}$$

where $Q = \text{diag}(Q_v, Q_\omega)$ and $R = \text{diag}(R_a, R_\delta)$.

Over the horizon, the tracking objective is written as

$$\sum_{t=0}^{T-1} \left\|D z_{t+1} - y_{t+1}^{\text{ref}}\right\|_Q^2. \tag{3.32}$$

The control smoothness term is

$$\sum_{t=0}^{T-1} \|u_t - u_{t-1}\|_R^2. \tag{3.33}$$

Hence, the total stage-wise cost is

$$J = \sum_{t=0}^{T-1} \left(\left\|D z_{t+1} - y_{t+1}^{\text{ref}}\right\|_Q^2 + \|u_t - u_{t-1}\|_R^2 \right). \tag{3.34}$$

To account for disturbances, we minimize the worst-case expected total cost over the ambiguity set $\mathcal{D}$. The distributionally robust cost function is defined as:

$$J_{\text{DR}} = \sup_{\mathbb{P}\in\mathcal{D}} \mathbb{E}_{\mathbb{P}} \left[\sum_{t=0}^{T-1} \left(\left\|D z_{t+1} - y_{t+1}^{\text{ref}}\right\|_Q^2 + \|u_t - u_{t-1}\|_R^2 \right) \right], \tag{3.35}$$

where the state and input sequences satisfy the lifted dynamics (3.23) under disturbance uncertainty.

Following the advisor's latest guidance, this manuscript version directly develops the tractable reformulation from per-step chance constraints and lifted dynamics.

4 Tractable Reformulation

To ensure safe obstacle avoidance under uncertain disturbances, we impose a per-step distributionally robust chance constraint as in (3.30). The objective simultaneously contains an expected cost over the same ambiguity set. These two components require different duality techniques: the worst-case expectation and the distributionally robust chance constraint.

4.1 Reformulation of the worst-case expected cost

Let $\hat{\mathbf{w}} = [\hat{w}_0^\top, \hat{w}_1^\top, \dots, \hat{w}_{T-1}^\top]^\top \in \mathbb{R}^{n_z T}$ denote the combined disturbance trajectory over the prediction horizon. For a fixed control sequence $\{u_t\}_{t=0}^{T-1}$ and initial lifted state $z_0 = \psi^e(s_{\text{current}})$, the total cost defined in (3.34) is

$$J(\hat{\mathbf{w}}) = \sum_{t=0}^{T-1} \left(\|Dz_{t+1} - y_{t+1}^{\text{ref}}\|_Q^2 + \|u_t - u_{t-1}\|_R^2 \right), \tag{4.1}$$

where the lifted states evolve according to (3.23):

$$z_{t+1} = Az_t + Bu_t + \hat{w}_t, \qquad t = 0, \dots, T-1. \tag{4.2}$$

Because the dynamics are affine, J is a convex quadratic function of $\hat{\mathbf{w}}$. The worst-case expected cost over the Wasserstein ambiguity set $\mathcal{D}$ (defined in (3.26) with radius θ and empirical distribution $\hat{\mathbb{P}}_N$ built from N historical disturbance trajectories $\{\hat{\mathbf{w}}^{(i)}\}_{i=1}^N$) is

$$R_{\text{exp}}(\{u_t\}) \;:=\; \sup_{\mathbb{P} \in \mathcal{D}} \mathbb{E}^{\mathbb{P}}[J(\hat{\mathbf{w}})]. \tag{4.3}$$

Theorem 1. *Assume that the disturbance $\hat{\mathbf{w}}$ is supported on a convex compact set $\Xi \subset \mathbb{R}^{n_z T}$ (physically justified, e.g., $\|\hat{\mathbf{w}}\|_\infty \le M$). Let κ be the Lipschitz constant of J with respect to the dual norm $\|\cdot\|_*$ on Ξ, i.e.,*

$$\kappa = \sup_{\hat{\mathbf{w}} \in \Xi} \|\nabla J(\hat{\mathbf{w}})\|_*. \tag{4.4}$$

Then the worst-case expected cost over the Wasserstein ambiguity set $\mathcal{D}$ satisfies

$$R_{exp}(\{u_t\}) \le \kappa\, \theta + \frac{1}{N} \sum_{i=1}^{N} J(\hat{\mathbf{w}}^{(i)}). \tag{4.5}$$

If $\Xi = \mathbb{R}^{n_z T}$, then

$$R_{exp}(\{u_t\}) = \kappa\, \theta + \frac{1}{N} \sum_{i=1}^{N} J(\hat{\mathbf{w}}^{(i)}). \tag{4.6}$$

i.e., the optimal control sequence that minimizes R_{exp} is identical to the one that minimizes the empirical average $\frac{1}{N}\sum_{i=1}^N J(\hat{\mathbf{w}}^{(i)})$.

Proof. By the dual representation of the Wasserstein metric (Esfahani & Kuhn [28, Eq. (12b)]), for any convex function ℓ on a convex set we have

$$\sup_{\mathbb{Q}\in\mathbb{B}_\theta(\hat{\mathbb{P}}_N)} \mathbb{E}^{\mathbb{Q}}[\ell(\xi)] = \inf_{\eta\geq 0}\left\{\eta\theta + \frac{1}{N}\sum_{i=1}^{N}\sup_{\xi\in\Xi}\big(\ell(\xi) - \eta\|\xi - \hat{\xi}_i\|\big)\right\}, \tag{4.7}$$

where Ξ is the support of ξ.
Case 1: Ξ compact. Since J is convex and Lipschitz on Ξ with constant κ, for any $\eta \geq \kappa$ we have

$$\sup_{\hat{\mathbf{w}}\in\Xi}\big(J(\hat{\mathbf{w}}) - \eta\|\hat{\mathbf{w}} - \hat{\mathbf{w}}^{(i)}\|\big) \leq J(\hat{\mathbf{w}}^{(i)}). \tag{4.8}$$

Choosing $\eta = \kappa$ in the dual representation gives an upper bound:

$$R_{\exp}(\{u_t\}) \leq \kappa\theta + \frac{1}{N}\sum_{i=1}^{N} J(\hat{\mathbf{w}}^{(i)}). \tag{4.9}$$

This proves the inequality in (4.5).
Case 2: $\Xi = \mathbb{R}^{n_z T}$. If $\eta < \kappa$, then because J has directional derivative κ in some direction and Ξ is unbounded, the supremum is $+\infty$. Therefore the infimum over η is restricted to $\eta \geq \kappa$. For $\eta \geq \kappa$, as before,

$$\sup_{\hat{\mathbf{w}}\in\mathbb{R}^{n_z T}}\big(J(\hat{\mathbf{w}}) - \eta\|\hat{\mathbf{w}} - \hat{\mathbf{w}}^{(i)}\|\big) = J(\hat{\mathbf{w}}^{(i)}), \tag{4.10}$$

with the supremum attained at $\hat{\mathbf{w}} = \hat{\mathbf{w}}^{(i)}$. Substituting into the dual representation yields

$$R_{\exp}(\{u_t\}) = \inf_{\eta\geq\kappa}\left\{\eta\theta + \frac{1}{N}\sum_{i=1}^{N} J(\hat{\mathbf{w}}^{(i)})\right\} = \kappa\theta + \frac{1}{N}\sum_{i=1}^{N} J(\hat{\mathbf{w}}^{(i)}), \tag{4.11}$$

where the infimum is attained at $\eta = \kappa$. This proves the equality in (4.6). □

4.2 Exact reformulation of distributionally robust chance constraints

For each obstacle j and time step t, the safety constraint (3.30) is

$$\inf_{\mathbb{P}\in\mathcal{D}} \Pr_{\mathbb{P}}\big(H_{\text{obs},j}Dz_t - h_{\text{obs},j} \leq 0\big) \geq 1 - \epsilon, \qquad \epsilon \in (0,1). \tag{4.12}$$

Define the loss function

$$\ell_{j,t}(\hat{\mathbf{w}}) = \big(H_{\text{obs},j}Dz_t - h_{\text{obs},j}\big)_+ = \max\big\{0,\ H_{\text{obs},j}Dz_t - h_{\text{obs},j}\big\}. \tag{4.13}$$

Using the lifted dynamics (3.23), z_t is an affine function of $\hat{\mathbf{w}}$:

$$z_t = A^t z_0 + \sum_{k=0}^{t-1} A^{t-1-k} B u_k \; + \; \sum_{k=0}^{t-1} A^{t-1-k} \hat{w}_k. \tag{4.14}$$

Define the constant row vector (for fixed j, t)

$$c_{j,t}^\top \;:=\; H_{\mathrm{obs},j} D \big[A^{t-1},\; A^{t-2},\; \ldots,\; I,\; 0,\; \ldots,\; 0\big] \in \mathbb{R}^{1 \times n_z T}, \tag{4.15}$$

and the deterministic term

$$d_{j,t} \;:=\; h_{\mathrm{obs},j} - H_{\mathrm{obs},j} D \left(A^t z_0 + \sum_{k=0}^{t-1} A^{t-1-k} B u_k \right), \tag{4.16}$$

which depends linearly on the control sequence. Then

$$H_{\mathrm{obs},j} D z_t - h_{\mathrm{obs},j} = c_{j,t}^\top \hat{\mathbf{w}} - d_{j,t}, \tag{4.17}$$

and therefore

$$\ell_{j,t}(\hat{\mathbf{w}}) = \big(c_{j,t}^\top \hat{\mathbf{w}} - d_{j,t}\big)_+ = \max\{0,\; c_{j,t}^\top \hat{\mathbf{w}} - d_{j,t}\}. \tag{4.18}$$

Condition (4.12) is equivalent (see Lemma 1 in [32]) to

$$\sup_{\mathbb{P} \in \mathcal{D}} \mathbb{E}^{\mathbb{P}}[\ell_{j,t}(\hat{\mathbf{w}})] \le \epsilon. \tag{4.19}$$

Theorem 2. *Let $\ell(\xi) = \max_{k \le K} \ell_k(\xi)$ where each $-\ell_k$ is proper, convex, and lower semicontinuous, and let the support of ξ be $\mathbb{R}^m$ (convex and closed). Then for any $\theta \ge 0$,*

$$\sup_{\mathbb{Q} \in \mathbb{B}_\theta(\hat{\mathbb{P}}_N)} \mathbb{E}^{\mathbb{Q}}[\ell(\xi)] = \inf_{\lambda \ge 0,\; s_i,\; z_{ik},\; v_{ik}} \left\{ \lambda\theta + \frac{1}{N} \sum_{i=1}^{N} s_i \right\}, \tag{4.20}$$

subject to

$$[-\ell_k]^*(z_{ik} - v_{ik}) + \sigma_{\mathbb{R}^m}(v_{ik}) - \langle z_{ik}, \hat{\xi}_i \rangle \le s_i, \tag{4.21}$$

$$\|z_{ik}\|_* \le \lambda, \qquad \forall i, k. \tag{4.22}$$

In the special case of right-hand uncertainty with $\ell(\xi) = (c^\top \xi - d)_+$ (i.e., $K = 2$, $\ell_1 = 0$, $\ell_2 = c^\top \xi - d$), this reduces to the linear constraints

$$\lambda\theta + \frac{1}{N} \sum_{i=1}^{N} s_i \le \epsilon, \tag{4.23}$$

$$s_i \ge c^\top \hat{\xi}_i - d + \lambda \|c\|_*, \quad s_i \ge 0, \quad \lambda \ge 0. \tag{4.24}$$

Proof. The first part is Theorem 4.2 of [31]. For the right-hand uncertainty case, set $\ell_1 = 0$, $\ell_2(\xi) = c^\top \xi - d$. Their conjugates are

$$[-\ell_1]^*(y) = \chi_{\{0\}}(y), \qquad [-\ell_2]^*(y) = \chi_{\{-c\}}(y) - d, \qquad \sigma_{\mathbb{R}^m}(v) = \chi_{\{0\}}(v). \tag{4.25}$$

Substituting into the general constraints (4.21)–(4.22) forces $v_{i1} = 0, z_{i1} = 0$ and $v_{i2} = 0, z_{i2} = -c$. The norm constraint gives $\|c\|_* \le \lambda$. The inequalities then become $0 \le s_i$ and $-d + c^\top \hat{\xi}_i \le s_i$. Minimising $\lambda\theta + \frac{1}{N}\sum s_i$ over $\lambda \ge \|c\|_*$ yields the optimal $\lambda = \|c\|_*$; keeping λ as a variable yields (4.23)–(4.24). □

Applying Theorem 2 to $\ell_{j,t}$ with $c = c_{j,t}$ and $d = d_{j,t}$, we obtain for each obstacle j and time t a set of linear constraints: there exist $\lambda_{j,t} \ge 0$ and $s_{j,t,i} \ge 0$ $(i = 1, \dots, N)$ such that

$$\lambda_{j,t}\theta + \frac{1}{N}\sum_{i=1}^{N} s_{j,t,i} \le \epsilon, \tag{4.26}$$

$$s_{j,t,i} \ge c_{j,t}^\top \hat{\mathbf{w}}^{(i)} - d_{j,t} + \lambda_{j,t}\|c_{j,t}\|_*. \tag{4.27}$$

These constraints are linear in the decision variables $(\lambda_{j,t}, s_{j,t,i}, u_t, z_t)$.

4.3 Distributionally robust MPC problem reformulation

Combining the cost reformulation and the chance-constraint reformulation (4.26)–(4.27), we obtain the following convex quadratic program:

$$\begin{aligned}
\min_{\{u_t, z_t\},\, \{\lambda_{j,t}\},\, \{s_{j,t,i}\}} \quad & \frac{1}{N}\sum_{i=1}^{N} J(\hat{\mathbf{w}}^{(i)}) \\
\text{s.t.} \quad & z_{t+1} = Az_t + Bu_t + \hat{w}_t, \quad u_t = Kz_t + v_t, \quad z_0 = \psi^e(s_{\text{current}}), \\
& v_{\min} \le Ez_t \le v_{\max}, \quad \forall t, \\
& \omega_{\min} \le Fz_t \le \omega_{\max}, \quad \forall t, \\
& a_{\min} \le a_t \le a_{\max}, \quad \forall t, \\
& |\delta_t| \le \delta_{\max}, \quad \forall t, \\
& |\delta_{t+1} - \delta_t| \le \Delta\delta_{\max}\Delta t, \quad \forall t, \\
& G_{\text{road},t} D z_t \le h_{\text{road},t}, \quad \forall t, \\
& \lambda_{j,t} \ge 0, \\
& \lambda_{j,t}\theta + \frac{1}{N}\sum_{i=1}^{N} s_{j,t,i} \le \epsilon, \quad \forall j, \forall t, \\
& s_{j,t,i} \ge 0, \quad s_{j,t,i} \ge c_{j,t}^\top \hat{\mathbf{w}}^{(i)} - d_{j,t} + \lambda_{j,t}\|c_{j,t}\|_*, \quad \forall j, t, i, .
\end{aligned} \tag{4.28}$$

The objective in (4.26) omits the constant term $\kappa\theta$ as it does not affect the optimal control; if a guaranteed upper bound on the out-of-sample cost is desired, $\kappa\theta$ can be added to the optimal value *a posteriori*. In our experiments, we impose no explicit constraints on Ξ and therefore assume $\Xi = \mathbb{R}^{n_z T}$. All constraints are linear in the decision variables, and the objective is quadratic because J is quadratic; hence (4.28) is a convex quadratic program that can be solved efficiently with standard solvers (e.g., OSQP, IPOPT, Gurobi). The reformulation is exact for the chance constraints under the assumptions of Theorem 2, and it provides a valid upper bound for the worst-case expected cost. Under the additional boundedness condition (or when the disturbance support is $\mathbb{R}^{n_z T}$ with finite κ), the bound becomes exact.

5 SIMULATION AND NUMERICAL RESULTS

This section evaluates the proposed Koopman-based distributionally robust MPC (K-DRMPC) in closed-loop simulation to assess whether distributionally robust constraints improve robustness under distribution shift while preserving nominal tracking performance. The evaluation examines complementary dimensions: (i) nominal performance parity under in-distribution disturbances relative to non-robust K-MPC, (ii) robustness gains under distribution shift in safety-related indicators such as success rate and lateral error percentiles, and (iii) the trade-off between safety, performance, and computational overhead as the parameters (θ, ϵ) are tuned. The evidence chain is organized around reproducibility protocol, failure mechanism analysis, distribution-shift robustness studies, and sensitivity studies.

5.1 Experimental Setup and Compared Methods

Experiments are performed on the Serpentine track for closed-loop tracking testing with boundary termination conditions. A run terminates early if the boundary condition is violated (e.g., $|e_y| \geq 10.0$ m), or when the maximum simulation steps are reached.

Compared methods. Two configurations are evaluated to isolate the effect of distributionally robust constraints:

i) K-MPC (Non-Robust): Standard Koopman MPC with deterministic obstacle avoidance constraints only. Constraint (2.15i) applies without the distributionally robust chance-constraint safety layer.
ii) K-DRMPC (Robust): Full distributionally robust formulation in Section 5, namely problem (4.28), incorporating the Wasserstein ambiguity set with radius $\theta > 0$ and risk level $\epsilon \in (0, 1)$.

Disturbance-shift protocol. In this study, the disturbance is modeled as an unknown additive term in the lifted dynamics, and no parametric family (e.g., Gaussian or Laplace) is assumed for the true distribution. Three disturbance settings are tested to create distributional shifts relative to the training data:

i) Setting A (Nominal, $\times 1.0$): The disturbance distribution matches the training set (empirical distribution $\hat{\mathbb{P}}_N$). K-MPC and K-DRMPC are expected to exhibit similar nominal performance.
ii) Setting B (Moderate Shift, $\times 1.5$): The disturbance magnitude increases moderately. The true distribution departs from $\hat{\mathbb{P}}_N$ but may still lie within a properly chosen Wasserstein ball. K-DRMPC is expected to maintain safety better than K-MPC.
iii) Setting C (Large Shift, $\times 2.0$): A severe out-of-distribution disturbance is applied. K-DRMPC with larger θ should retain stronger safety guarantees, while K-MPC is more likely to fail.

Vehicle model and sampling. The controlled plant is a kinematic bicycle model with state $s_t = [p_{x,t}, p_{y,t}, \psi_t, v_t, \omega_t]^\top$ and input $u_t = [a_t, \delta_t]^\top$. The sampling time is $\Delta t = 0.1$ s and the wheelbase is $L = 2.6$ m. The MPC prediction horizon is set to $T = 20$ steps (2.0 s). The Koopman latent dimension is $n_z = 32$.

Controller and solver settings. Table 5.1 summarizes the key parameters used in simulation and optimization.

Table 5.1 Simulation and optimization parameters (shared by K-MPC and K-DRMPC).

Item	**Symbol**	**Value**
Sampling time	Δt	0.1 s
Wheelbase	L	2.6 m
Prediction horizon	T	20 steps
Latent dimension	n_z	32
Speed limits	$[v_{\min}, v_{\max}]$	[0.0, 4.0] m/s
Acceleration limits	$[a_{\min}, a_{\max}]$	[-0.6, 0.5] m/s$\hat{2}$
Steering limit	$\delta_{\max}$	12.0 deg
Steering-rate limit	$\dot{\delta}_{\max}$	2.00 rad/s
Geometric safety margin	η	0.50 m
Solver	–	IPOPT
Non-Robust (K-MPC) specific:		
Obstacle avoidance	(2.15i)	Deterministic inequality
Robust (K-DRMPC) specific:		
Wasserstein radius	θ	0.100
Chance-constraint risk level	ϵ	0.100
Obstacle avoidance	(4.28)	Distributionally robust chance constraints

Metrics and statistical protocol. Table 5.2 summarizes all reported metrics and their statistical definitions.

For obstacle-avoidance evaluation, the signed safety margin at time step t is defined as

$$m_t := h_{\text{obs},j} - H_{\text{obs},j} D z_t, \tag{5.1}$$

where $m_t > 0$ indicates safe separation, $m_t = 0$ indicates boundary contact, and $m_t < 0$ indicates violation. At the trajectory level, we report the minimum safety margin

$$m_{\min} := \min_t m_t, \tag{5.2}$$

Table 5.2 Metrics and statistical protocol used in Section 6.

Category	Metric	Definition / Statistic
Tracking	Lateral error $\lvert e_y \rvert$	Mean, 95th percentile, and maximum over closed-loop steps.
Tracking	Heading error $\lvert e_\psi \rvert$	Mean, 95th percentile, and maximum over closed-loop steps.
Safety	Completion status	Binary indicator (completed / terminated).
Safety	Minimum safety margin $m_{\min}$	With signed margin $m_t = h_{\mathrm{obs},j} - H_{\mathrm{obs},j} Dz_t$, report $m_{\min} = \min_t m_t$ (larger is safer; negative indicates violation).
Safety	Safety-margin quantile m_{P05}	5th percentile of $\{m_t\}$ over closed-loop steps, used to reduce sensitivity to single-step outliers.
Safety	Crash step / crash time	First step and corresponding time when boundary/safety violation occurs.
Safety	Violation rate	Ratio of violating steps to total simulated steps.
Computational	Solve time per MPC step	Mean solve time (ms) over all optimization calls.
Computational	Solve-time percentiles	P50, P95, and P99 of solve time (ms).
Stability	Longest out-of-band interval	Maximum consecutive duration with $\lvert e_y \rvert > 2$ m (in steps or seconds).

as a single scalar safety indicator, together with the 5th-percentile statistic m_{P05} to reduce sensitivity to single-step outliers.

When repeated experiments are available, all scalar metrics should additionally report mean±std and confidence intervals. In the following subsection, safety-margin statistics are reported together with completion status to make the safety-performance trade-off explicit.

5.2 Simulation Results and Sensitivity Analysis

Figure 5.1 first provides a geometric view of the closed-loop behavior in the primary scenario. The robust controller follows a smoother path around obstacle clusters and maintains larger boundary clearance at turning segments, while the non-robust controller shows visibly larger path excursions in high-curvature regions. This qualitative contrast is consistent with the quantitative evidence reported later in Table 5.3 and supports the interpretation that robustness mainly improves feasibility and tail behavior under distribution shift. More specifically, the primary experimental differences can be summarized along three dimensions. First, reliability is the dominant separator: Settings A and B do not distinguish the controllers in terms of completion, but under the severe shift of Setting C the comparison changes from nominal parity to a clear feasibility gap, indicating that the main value of the robust formulation lies in preserving task completion under mismatch rather than merely improving benign-regime averages. Second, the robust advantage is concentrated in the error tails. Even when mean tracking differences remain moderate, K-DRMPC consistently reduces large excursions, especially in the P95 lateral and heading errors, which explains why the geometric trajectories appear visibly smoother in the most demanding regions. Third,

this robustness gain is not free: it is achieved at a moderate but stable computational surcharge, meaning that the method should be interpreted as exchanging additional solve time for substantially improved mismatch tolerance. For safety-critical deployment, this trade-off is favorable because failure prevention under distribution shift is more decision-relevant than small nominal-speed differences.

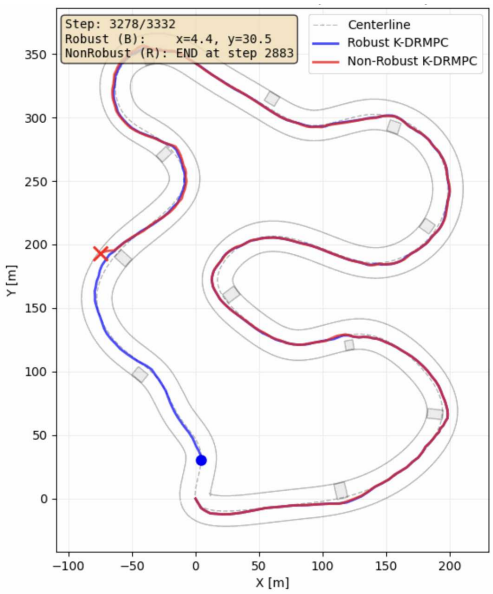


Fig. 5.1 Closed-loop trajectory on the primary scenario. The dashed curve is the reference centerline of a preplanned serpentine track, and about 10 gray rectangles denote the fixed static obstacles used throughout the experiments.

Table 5.3 reports the core closed-loop outcomes for K-MPC and K-DRMPC across three disturbance settings (A: nominal, B: moderate shift, C: severe shift), enabling direct comparison of robustness performance under distributional changes. For Setting C, note that early termination appears in repeated runs (all K-MPC runs and part of K-DRMPC runs); therefore, trajectory-averaged metrics are computed over truncated horizons and are not strictly comparable to full-horizon runs.

Parameter sensitivity analysis: To characterize the internal robustness trade-off of K-DRMPC, we sweep the Wasserstein radius $\theta \in \{0.0, 0.5, 1.0, 1.5, 2.0\}$ and the chance constraint risk level $\epsilon \in \{0.01, 0.05, 0.10\}$ under Setting B. Figure 5.3 summarizes the resulting sensitivity using the lower-tail safety-margin statistic m_{P05}.

During the experiments, we made three main observations. First, all 15 parameter pairs remained feasible, indicating that the controller retains a broad operating region in this moderate-shift regime. Second, m_{P05} is much more sensitive to ϵ than to θ: decreasing ϵ from 0.10 to 0.01 consistently increases the lower-tail safety margin, so the chance constraint level is the primary conservatism knob. Third, the effect of θ is comparatively weak and non-monotone. For example, when $\epsilon = 0.10$, the heatmap is nearly flat except for a local dip around $\theta = 1.0$; when $\epsilon = 0.01$, the margin

Table 5.3 Closed-loop performance comparison across Settings A, B, C ($\theta = 0.10$, $\epsilon = 0.10$).

Metric	Setting A (Nominal)		Setting B (Moderate Shift)		Setting C (Severe Shift)	
	K-MPC	**K-DRMPC**	**K-MPC**	**K-DRMPC**	**K-MPC**	**K-DRMPC**
Trajectory Completion						
Steps	3378.7	3371.0	3401.0	3387.3	1198.3	2397.0
Task time (s)	337.9	337.1	340.1	338.7	119.8	239.7
Crash (Yes/No)	No	No	No	No	Yes (3/3)	Mixed (1/3)
Safety-Margin Statistics						
$m_{\min}$ (m)	-2.793	-2.681	-3.426	-3.355	-1.016	-2.319
m_{P05} (m)	-0.361	-0.506	-0.503	-0.633	0.852	0.062
Lateral Error						
Mean $\lvert e_y \rvert$ (m)	1.374	1.300	1.829	1.652	2.041	2.080
P95 $\lvert e_y \rvert$ (m)	3.885	3.598	4.920	4.320	6.232	5.866
Heading Error						
Mean $\lvert e_\psi \rvert$ (deg)	3.048	2.772	3.923	3.531	5.685	4.333
P95 $\lvert e_\psi \rvert$ (deg)	9.825	9.252	11.986	11.052	25.534	13.127
Velocity & Computational						
Mean speed (m/s)	4.313	4.311	4.313	4.307	4.942	4.502
Solve mean (ms)	1373.9	1421.1	1358.0	1394.6	1345.7	1392.4
Solve P95 (ms)	1493.2	1553.2	1479.6	1523.1	1507.0	1549.2



Fig. 5.2 Aggregated time-series comparison between K-MPC and K-DRMPC for Setting C (severe disturbance, $\sigma = 1.0$). Curves show run-wise means from the full log batch (three runs per method), with shaded bands indicating dispersion. The safe margin panels (top two) reveal the critical advantage: K-MPC's obstacle and track margins collapse to zero around $t = 280$ s, causing mission failure, whereas K-DRMPC maintains positive margins throughout. The remaining four panels report $|e_y|$, $|e_\psi|$, speed profile v, and per-step MPC solve time.

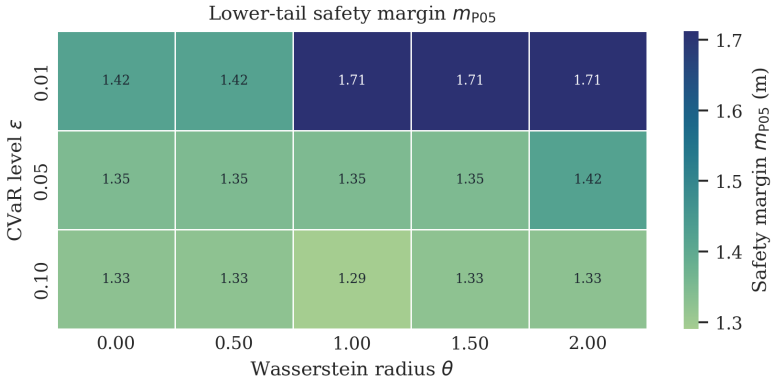


Fig. 5.3 Heatmap of the lower-tail safety margin m_{P05} over the full (θ, ϵ) grid under Setting B ($\sigma = 0.75$). Each cell reports the realized m_{P05} for one parameter pair; the color field visualizes how the robustness configuration changes the lower-tail clearance.

increases once θ reaches 1.0 and then saturates. Overall, ϵ is the more effective tuning parameter for enlarging realized safety margins, while θ mainly provides a secondary robustness-shaping effect over the tested range.

In addition, the comparison between robust and non-robust control yields several clear observations. Under the nominal condition (Setting A), both controllers completed all three runs. K-DRMPC achieved slightly better tracking tails: the mean lateral error dropped from 1.374 m to 1.300 m, and the 95th percentile from 3.885 m to 3.598 m. Heading errors also improved, with the mean and 95th percentile falling from 3.048° and 9.825° to 2.772° and 9.252°, respectively. This came at a computational cost: the mean solve time increased from 1373.9 ms to 1421.1 ms.

Under the moderate shift (Setting B), both methods still finished all runs, but the robust controller showed larger gains. Its lateral error mean and 95th percentile improved to 1.652 m and 4.320 m, compared to 1.829 m and 4.920 m for K-MPC. The fraction of steps where lateral error stayed within 2 m rose from 66.14% to 69.91%. This suggests that the DRO constraints already help under distribution shifts even before any hard failures appear.

The severe shift (Setting C) produced the clearest difference. K-MPC failed in all three seeds, whereas K-DRMPC completed two of three runs. The robust controller also cut the extreme error tails markedly: the 95th percentile of lateral error went from 6.232 m to 5.866 m, and that of heading error from 25.534° to 13.127°. However, mean error numbers are less comparable here because early termination truncates the trajectories and changes the sample composition. Safety margins such as $m_{\min}$ and m_{P05} should always be interpreted together with completion status. In Setting C, the early termination of non-robust runs can artificially inflate quantiles (e.g., giving a positive m_{P05} despite zero completed runs). Hence, under a severe shift, the completion rate is the primary safety indicator.

From a computational and structural perspective, the Koopman-based formulation strikes a practical balance between model expressiveness and online tractability. Lifting the nonlinear dynamics into a high-dimensional latent space allows MPC to apply

linear control theory while preserving global validity. Early-step control decisions and steering-rate constraints have a strong impact on safety and performance. An incorrect steering sign at the beginning can cause an irreversible boundary violation; raising the steering-rate limit could mitigate this. The runtime of K-DRMPC scales with the number of historical disturbance samples (due to the chance constraints, which introduce slack variables $s_{t,i}$) and with the prediction horizon T. Across Settings A–C, the robust overhead in mean solve time ranged from about +36.6 ms to +47.1 ms, and the 95th percentile overhead from +42.2 ms to +60.0 ms.

6 CONCLUSION

This paper has presented a Koopman-based distributionally robust MPC framework (K-DRMPC) for autonomous vehicle obstacle avoidance under nonlinear dynamics and unknown disturbance distributions. The proposed method integrates three key ideas: global linearization via deep Koopman lifting, a Wasserstein ambiguity set to capture distributional uncertainty, and a tractable reformulation of the resulting chance constraints using conditional value-at-risk. The main contributions are threefold. First, a control-oriented Koopman formulation lifts the nonlinear vehicle dynamics into a latent space where linear MPC techniques apply, thereby achieving global validity beyond local equilibrium points. Second, a distributionally robust safety framework explicitly accounts for model mismatch and external disturbances through a Wasserstein ambiguity set, balancing performance and safety under distributional shifts. Third, averaged simulation results under three progressively increasing disturbance levels demonstrate that K-DRMPC preserves nominal performance, improves tail tracking quality under moderate disturbances, and achieves substantially better feasibility when disturbances become severe—completing two out of three runs, compared to zero completions by the non-robust K-MPC.

The essential difference between the robust and non-robust controllers is not their nominal accuracy but their reliability as the disturbance magnitude grows. Under nominal conditions, both controllers complete all runs with comparable tracking quality, confirming that robustness does not meaningfully impair normal operation. When the disturbance level increases moderately, K-DRMPC already exhibits a clear advantage in tail errors while still completing every run. Under severe disturbances, the gap becomes decisive: K-MPC fails in all runs, whereas K-DRMPC succeeds in two out of three. This outcome indicates that distributional robustness primarily serves as a feasibility-preserving mechanism, though it incurs a moderate computational overhead. For safety-critical applications, completion under large disturbances must therefore take precedence over single-metric averages, and margin quantiles should always be interpreted alongside truncation and crash information.

Funding. This work was supported by the National Natural Science Foundation of China (No. 52432012).

Data availability. The source code and the simulation data necessary to reproduce the findings of this study are publicly available in the GitHub repository https://github.com/xuenanhkust163/Koopman-DRMPC-ObstacleAvoidance.

References

[1] B. O. Koopman, "Hamiltonian systems and transformation in Hilbert space," *Proc. Natl. Acad. Sci. USA*, vol. 17, no. 5, pp. 315–318, 1931.

[2] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, "A data–driven approximation of the Koopman operator: Extending dynamic mode decomposition," *J. Nonlinear Sci.*, vol. 25, no. 6, pp. 1307–1346, 2015.

[3] B. Lusch, J. N. Kutz, and S. L. Brunton, "Deep learning for universal linear embeddings of nonlinear dynamics," *Nat. Commun.*, vol. 9, no. 1, p. 4950, 2018.

[4] H. Shi and M. Q.-H. Meng, "Deep Koopman operator with control for nonlinear systems," *IEEE Robot. Autom. Lett.*, vol. 7, no. 3, pp. 7700–7707, 2022.

[5] C. Gu, T. Zhou, and C. Wu, "Deep Koopman traffic modeling for freeway ramp metering," *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 6, pp. 6001–6013, 2023.

[6] H. Zheng, Y. Li, L. Zheng, and E. Hashemi, "Koopman-based hybrid modeling and zonotopic tube robust MPC for motion control of automated vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 25, no. 10, pp. 13598–13612, 2024.

[7] X. Zhang, W. Pan, R. Scattolini, S. Yu, and X. Xu, "Robust tube-based model predictive control with Koopman operators," *Automatica*, vol. 137, p. 110114, 2022.

[8] X. Xiao, Y. Zhang, and J. Wang, "Deep neural networks with Koopman operators for modeling and control of autonomous vehicles," *IEEE Trans. Intell. Veh.*, vol. 7, no. 3, pp. 654–665, 2022.

[9] P. M. Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations," *Math. Program.*, vol. 171, no. 1, pp. 115–166, 2018.

[10] B. Li, Y. Tan, A.-G. Wu, and G.-R. Duan, "A distributionally robust optimization based method for stochastic model predictive control," *IEEE Trans. Autom. Control*, vol. 67, no. 11, pp. 5762–5777, 2021.

[11] A. Hakobyan and I. Yang, "Wasserstein distributionally robust motion control for collision avoidance using conditional value-at-risk," *IEEE Trans. Robot.*, vol. 39, no. 4, pp. 2690–2707, 2023.

[12] L. Chen, Y. Liu, and Z. Li, "Distributionally robust motion planning for autonomous driving in uncertain environments," *IEEE Trans. Intell. Transp. Syst.*, vol. 25, no. 1, pp. 567–578, 2024.

[13] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6,

pp. 789–814, 2000.

[14] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, 2nd ed. Santa Barbara, CA, USA: Nob Hill Publishing, 2017.

[15] X. Zhang, X. Xu, and S. Yu, "Koopman model predictive control for nonlinear systems," *Automatica*, vol. 132, p. 109800, 2021.

[16] B. Li, T. Guan, L. Dai, and G.-R. Duan, "Distributionally robust model predictive control with output feedback," *IEEE Trans. Autom. Control*, vol. 68, no. 9, pp. 5443–5458, 2023.

[17] Y. Wang, X. Zhang, and J. Wang, "Data-driven predictive control for autonomous vehicles using Koopman operator," *IEEE Trans. Intell. Transp. Syst.*, vol. 25, no. 2, pp. 1234–1245, 2024.

[18] J. Liu, X. Zhao, B. Zhu, J. Zhao, and W. Deng, "Model predictive control of vehicle dynamics based on the Koopman operator with extended dynamic mode decomposition," *IEEE Access*, vol. 12, pp. 57557–57568, 2024.

[19] M. Werling, J. Ziegler, S. Kammel, and S. Thrun, "Optimal trajectory generation for dynamic street scenarios via convex optimization," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Anchorage, AK, USA, May 2010, pp. 5016–5022.

[20] M. Werling, G. Hoffmann, B. Spillmann, and S. Thrun, "Online trajectory planning and tracking for autonomous off-road vehicles," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Anchorage, AK, USA, May 2010, pp. 5084–5089.

[21] M. Abtahi, F. Motallebi Araghi, N. Mojahed, and S. Nazari, "Deep bilinear Koopman model for real-time vehicle control in Frenet frame," *IEEE Trans. Intell. Veh.*, under review, 2025. [Online]. Available: arXiv:2507.12578.

[22] H. Chen and C. Lv, "Incorporating ESO into deep Koopman operator modelling for control of autonomous vehicles," arXiv preprint arXiv:2405.10157, 2024.

[23] S. Fang, X. Li, C. Wu, and K. Yu, "DRCC-LPVMPC: Robust data-driven control for autonomous driving and obstacle avoidance," arXiv preprint arXiv:2603.14408, 2026.

[24] A. da C. Vangasse, E. J. R. Freitas, G. V. Raffo, and L. C. A. Pimenta, "Safe Navigation on Path-Following Tasks: A Study of MPC-based Collision Avoidance Schemes in Distributed Robot Systems," *J. Intell. Robot. Syst.*, vol. 110, no. 4, 2024, doi: 10.1007/s10846-024-02202-3.

[25] J. Chen, H. Ebel, and P. Eberhard, "Practical Insights on Data-Based Robot Control: A Comparative Analysis of Data-Enabled Predictive Control and Model-Based Predictive Control," *J. Intell. Robot. Syst.*, vol. 112, no. 2, 2026, doi: 10.1007/s10846-026-02386-w.

[26] T. Weiss and M. Behl, "BARTe: Composite Bezier Curves for Racing Trajectory Estimation," *J. Intell. Robot. Syst.*, vol. 112, no. 2, 2026, doi: 10.1007/s10846-026-02368-y.

[27] M. Nezami, D. S. Karachalios, G. Schildbach, and H. S. Abbas, "Obstacle Avoidance of Autonomous Vehicles: An LPVMPC with Scheduling Trust Region," arXiv preprint arXiv:2405.02030, 2024.

[28] E. L. Zhu and F. Borrelli, "A sequential quadratic programming approach to the solution of open-loop generalized Nash equilibria for autonomous racing," arXiv preprint arXiv:2404.00186, 2024.

[29] H. Xue, E. L. Zhu, J. M. Dolan, and F. Borrelli, "Learning model predictive control with error dynamics regression for autonomous racing," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Yokohama, Japan, May 2024, pp. 4614–4620.

[30] L. Numerow, A. Zanelli, A. Carron, and M. N. Zeilinger, "Inherently robust suboptimal MPC for autonomous racing with anytime feasible SQP," arXiv preprint arXiv:2401.02194, 2024.

[31] Peyman Mohajerin Esfahani and Daniel Kuhn, *Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations*, Mathematical Programming, vol. 171, no. 1, pp. 115–166, 2018.

[32] Weijun Xie, *On distributionally robust chance constrained programs with Wasserstein distance*, Mathematical Programming, vol. 186, pp. 115–155, 2021. (Published online 13 November 2019)