
Nonlinear Wasserstein Distributionally Robust Optimal Control

Zhengang Zhong¹ Jia-Jie Zhu²

Abstract

This paper presents a novel approach to addressing the distributionally robust nonlinear model predictive control (DRNMPC) problem. Current literature primarily focuses on the static Wasserstein distributionally robust optimal control problem with a prespecified ambiguity set of uncertain system states. Although a few studies have tackled the dynamic setting, a practical algorithm remains elusive. To bridge this gap, we introduce a DRNMPC scheme that dynamically controls the propagation of ambiguity, based on the constrained iterative linear quadratic regulator. The theoretical results are also provided to characterize the stochastic error reachable sets under ambiguity. We evaluate the effectiveness of our proposed iterative DRMPC algorithm by comparing the closed-loop performance of feedback and open-loop on a mass-spring system, and demonstrate in numerical experiments that our algorithm controls the propagated Wasserstein ambiguity.

1. Introduction

Model predictive control (MPC) repeatedly solves optimization problems online based on a system model and prescribed constraints to determine optimal control actions (Mayne et al., 2000). However, the closed-loop performance of MPC designed based on the nominal system model could be severely deteriorated when the real system suffers from uncertainty (Cannon et al., 2009). To effectively develop control methods addressing the detrimental effect of uncertainty, two classes of MPC that explicitly take the uncertainty into account have emerged: stochastic MPC (SMPC) and robust MPC (RMPC). RMPC determines the

optimal control actions under the worst-case scenario within a pre-specified deterministic uncertainty set (Mayne et al., 2005), whereas SMPC assumes or estimates the distribution of the uncertainty and selects the best control action for an objective function under soft constraints (Mayne, 2014). However, the performance of RMPC might be over-conservative as low-probability uncertainty is also taken into account, whereas the actual performance of SMPC could significantly deviate from the designed one due to the distribution discrepancy between the true distribution and the nominal distribution used in the controller design (Heirung et al., 2018).

For the purpose of addressing the challenges mentioned above - conservativeness or misspecified nominal distribution - we consider a data-driven distributionally robust nonlinear MPC (DRNMPC) problem using the Wasserstein metric. In the construction of this controller, instead of knowing the probability distribution of disturbances exactly, only samples of the disturbance realizations are required to construct the Wasserstein ambiguity set. The ambiguity set includes the empirical distribution of disturbance samples at its center and all distributions within a certain Wasserstein distance. Control actions are determined based on the worst-case distribution from this set, considering distributional robustness.

1.1. Related work

Recently, distributionally robust control using the Wasserstein ambiguity garners a lot of interest and attention. For the purpose of state constraint satisfaction, the recent papers (Mark & Liu, 2020; Zhong et al., 2021; Coulson et al., 2021; Micheli et al., 2022; Fochesato & Lygeros, 2022) consider such a distributionally robust MPC problem with respect to the Wasserstein ambiguity set defined on the product probability space for linear systems, wherein the center of the ambiguity set is determined based i.i.d. samples of disturbance sequences. Both (Hakobyan & Yang, 2022) and (Yang, 2020) consider the distributionally robust control problem as a two-player zero-sum game without state constraints and solve the problem via dynamic programming with a relaxed formulation using a Wasserstein penalty. For nonlinear systems, (Zolanvari & Cherukuri, 2022) considers deterministic systems with disturbed constraints, and (Zhong et al., 2023) solves distributionally robust MPC

¹Sargent Centre for Process Systems Engineering, Imperial College London, London, UK ²Weierstrass Institute for Applied Analysis and Stochastics, Berlin, Germany. Correspondence to: Zhengang Zhong <z.zhong20@imperial.ac.uk>, Jia-Jie Zhu <zhu@wias-berlin.de>.




Figure 1. Closed-loop error dynamics (as in tube MPC) of 30 realizations with the feedback gains and nominal inputs solved by 1. Red: Our method. Blue: Fixed feedback gain. Black: Zero feedback gain, i.e. open-loop. The arrow indicates the error vector between two consecutive sampling times, i.e. the tail indicates the accumulated error of all the previous steps and the head indicates the accumulated error including the error from the current step. See Section 4.4 for the theoretical characterization.

for nonlinear systems with additive disturbances via iterative linearization. However, in (Zhong et al., 2023), the propagation of the ambiguity set is not considered and the feedback gain is static for each sampling time. After the initial submission of this manuscript, we were brought to the awareness of a recent preprint (Aolaritei et al., 2023), which is the closest to our work. While both works consider the dynamic setting in terms of Wasserstein ambiguity, we directly formulate the DRO problem based on the disturbance ambiguity dynamically, which is mathematically equivalent to their propagation to the state distributions while enjoying simpler forms; cf. (18), (27). Furthermore, compared with (Aolaritei et al., 2023), this paper solves DRNMPC based on iterative LQR, whereas they considered linear systems.

In this work, we consider Wasserstein distributionally robust MPC for nonlinear systems with additive disturbances. Instead of constructing the Wasserstein ambiguity set for disturbance sequences, we consider the Wasserstein ambiguity set of disturbance for the single-step dynamics and propagate the Wasserstein ambiguity sets within the prediction horizon. Also, instead of using a relaxed formulation, we solve the original DRNMPC problem via an iterative method with the help of Riccati recursion. We will show that our method could dynamically control the propagation of the Wasserstein ambiguity sets and hence guarantee a non-conservative closed-loop performance.

1.2. Contribution

This paper makes the following main contributions. 1) We solve a Wasserstein distributionally robust nonlinear model predictive control (DRNMPC) problem for nonlinear systems (5). To the best of our knowledge, this is the only work

that does not assume a priori Wasserstein ambiguity sets of the state distributions for nonlinear systems. 2) We present an iterative-linearized DRMPC scheme that uses feedback to dynamically control the propagation of Wasserstein ambiguity sets, whereas open-loop control fails to do so. The derivation of such an approach is summarized in Proposition 4.5 and the corresponding algorithm is introduced in Algorithm 1. This approach is a significant improvement over the existing literature, as previous research only addressed static problems with fixed ambiguity sets or dynamic problems without a practical algorithm. To the authors' best knowledge, our proposed algorithm is the first to provide a practical and efficient method for controlling the propagation of Wasserstein ambiguity sets in nonlinear dynamics. 3) We analytically characterize the Wasserstein distributional reachable set under dynamic propagation in our algorithm in Proposition 4.8. 4) We visualize the closed-loop performance of the proposed approach via an error diagram in Fig 1. We observe that our method effectively controls the propagation of the ambiguity sets.

The rest of the paper is organized as follows. In Section II, we introduce the control problem and the preliminary DRNMPC. Section III describes the Wasserstein ambiguity set applied to this work. In Section IV, we introduce the propagation of Wasserstein ambiguity sets and the corresponding algorithm dynamically controls the propagation. Also, we analyze the reachable sets of dynamic Wasserstein ambiguity and linearization error. In Section V, we provide a numerical experiment of a mass-spring system to demonstrate our method and comparison results.

2. Problem statement

2.1. Notations

We use x_k for the measured state at time k and $x_{i|k}$ for the state predicted i steps ahead at time k . $[A]_j$ and $[a]_j$ denote the j -th row and entry of the matrix A and vector a , respectively. Similarly, we denote the element of i -th row and j -th column in the matrix A as $[A]_{ij}$. We define the notation $[A]_{i:j}$ for the i -th to j -th row in the matrix A . The set $\mathbb{N}_{>0}$ denotes the positive integers and $\mathbb{N}_{\geq 0} = \{0\} \cup \mathbb{N}_{>0}$. The set \mathbb{N}_1^N denotes the set of integers from 1 to N . $\mathcal{M}(\Xi)$ defines the space of all probability distributions supported on Ξ with finite first moments. $(\cdot)^{(i)}$ denotes the i -th sample from the training set. The sequence of length N of vectors $v_{0|k}, \dots, v_{N-1|k}$ is denoted by $\mathbf{v}_{N|k}$. γ_{ij} denotes the element of a 2-D tensor, such that this element is the i -th, j -th element along the first and second axis, respectively. Similar for 1-D λ_i . Let $\mathbb{B}_{\infty}^{n_x} := \{d \in \mathbb{R}^{n_x} \mid \|d\|_{\infty} \leq 1\}$ denote the unit ball. Let $\mathbb{P}^{\otimes i} := \mathbb{P}_0 \times \dots \times \mathbb{P}_{i-1}$ denote the product distribution.

2.2. System dynamics, constraints and objective

We consider the nonlinear time-invariant dynamical system with additive disturbance

$$x_{k+1} = f_d(x_k, u_k) + w_k, \quad k \in \mathbb{N}_{\geq 0}, \quad (1)$$

where $f_d : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ is a discrete-time nonlinear dynamics, k is the discrete sampling time, the state $x_k \in \mathbb{R}^{n_x}$, the control $u_k \in \mathbb{R}^{n_u}$, and the additive disturbance $w_k \in \mathbb{R}^{n_x}$. Each disturbance w_k of the disturbance sequence $\{w_k\}_{k \in \mathbb{N}_{\geq 0}}$ is assumed to be a realization of the corresponding random variable (r.v.) W_k from the random process $\{W_k\}_{k \in \mathbb{N}_{\geq 0}}$ satisfying the following assumption.

Assumption 2.1 (Bounded i.i.d Random Disturbance). All random variables $W_k \sim \mathbb{P}_w$ for $k \in \mathbb{N}_{\geq 0}$ from the family of random variables $\{W_k\}_{k \in \mathbb{N}_{\geq 0}}$ are assumed to be zero-mean and independent and identically distributed (i.i.d) with an unknown probability distribution \mathbb{P}_w and a known polyhedral support $\mathbb{W}_w \triangleq \{w \mid H_w w \leq h_w\}$.

For any given state measurement x_k at the sample time k , the predicted system states within the prediction horizon N are described as

$$x_{i+1|k} = f_d(x_{i|k}, u_{i|k}) + W_{i|k}, \quad x_{0|k} = x_k,$$

where $x_{i|k}$, $u_{i|k}$, and $W_{i|k} := W_{k+i}$ are all random variables.

We further introduce the nonlinear dynamics $f_{d,i} : \mathbb{R}^{n_x} \times \underbrace{\mathbb{R}^{n_u} \times \dots \times \mathbb{R}^{n_u}}_{i \text{ times}} \times \underbrace{\mathbb{R}^{n_x} \times \dots \times \mathbb{R}^{n_x}}_{i \text{ times}} \rightarrow \mathbb{R}^{n_x}$ for the predicted state $x_{i|k}$ with $i \geq 1$ dependent on the measurement x_k , input sequence $u_{0|k}, \dots, u_{i-1|k}$, and disturbance sequence $W_{0|k}, \dots, W_{i-1|k}$

$$\begin{aligned} x_{i|k} &= f_{d,i}(x_k, u_{0|k}, \dots, u_{i-1|k}, W_{0|k}, \dots, W_{i-1|k}) \\ &:= f_d(f_d(\dots f_d(x_k, u_{0|k}) + W_{0|k} \dots), u_{i-1|k}) \\ &\quad + W_{i-1|k}. \end{aligned} \quad (2)$$

To highlight that the predicted state is dependent on the disturbance sequence, we use a slight abuse of notation and denote

$$\begin{aligned} f_{d,i}(x_k, u_{0|k}, \dots, u_{i-1|k}, W_{0|k}, \dots, W_{i-1|k}) &:= \\ x_{i|k}(W_{0|k}, \dots, W_{i-1|k}). \end{aligned}$$

For any nonlinear system, we consider distributionally robust constraints with ambiguity set propagation to the states

$$\begin{aligned} \sup_{\mathbb{P}_m \in \mathcal{P}_{k+m}, m=0, \dots, i-1} \mathbb{E}_{\mathbb{P}^{\otimes i}} \{ [F]_n x_{i|k}(W_{0|k}, \dots, W_{i-1|k}) \} \\ \leq [f]_n, \end{aligned} \quad (3)$$

where $W_{m|k} \sim \mathbb{P}_m$ is the disturbance variable, $\mathbb{P}^{\otimes i} := \mathbb{P}_0 \times \dots \times \mathbb{P}_{i-1}$, $k \in \mathbb{N}_{\geq 0}$, $n \in \mathbb{N}_1^{n_F}$, $i \in \mathbb{N}_1^N$, $F \in$

$\mathbb{R}^{n_F \times n_x}$, $f \in \mathbb{R}^{n_F}$. For each additive disturbance $W_{i|k}$ within the prediction horizon, we centered an ambiguity set \mathcal{P}_{k+i} as the Wasserstein ball around the empirical distribution $\hat{\mathbb{P}}_{k+i} := \frac{1}{M} \sum_{l=1}^M \delta_{\hat{w}_{i|k}^{(l)}}$. Due to the i.i.d assumption, the realization of additive disturbance is time-independent; hence, we will denote the ambiguity set as \mathcal{P} and the corresponding empirical distribution as $\hat{\mathbb{P}}$ without explicitly indicating the predicted step i .

Remark 2.2. Through our formulation of the worst-case distributionally robust state constraints, the predicted states are affected by the accumulated error of the worst-cast distributions from each previous step within the prediction horizon. Hence, the control actions will be determined with an explicit consideration of the propagated effect of the worst distribution at each step of prediction. More details will be introduced in Section 4.

Without loss of generality, we characterize the control target as tracking the equilibrium point, which we assume to be the origin of the coordinate system, from an initial state while satisfying the prespecified constraints. The control objective is hence defined as the minimization of the expected value with the reference trajectory uniformly equal to zero

$$\mathbb{E}_{\mathbb{P}} \left\{ \sum_{i=0}^{N-1} (\|x_{i|k}\|_Q^2 + \|u_{i|k}\|_R^2) + \|x_{N|k}\|_{Q_f}^2 \right\}. \quad (4)$$

Here $Q, Q_f \in \mathbb{R}^{n_x \times n_x}$ and $\mathbb{R}^{n_u \times n_u}$ are penalty matrices for the quadratic stage costs. The corresponding optimization problem of DRNMPC for nonlinear systems is defined as

Problem 1.

$$\begin{aligned} \min_{\mathbf{u}} \quad & \mathbb{E}_{\mathbb{P}} \left\{ \sum_{i=0}^{N-1} (\|x_{i|k}\|_Q^2 + \|u_{i|k}\|_R^2) + \|x_{N|k}\|_{Q_f}^2 \right\} \\ \text{s.t.} \quad & x_{0|k} = x_k \\ & x_{i+1|k} = f_d(x_{i|k}, u_{i|k}) + W_{k+i} \\ & \sup_{\mathbb{P}_m \in \mathcal{P}_{k+m}, m=0, \dots, i-1} \mathbb{E}_{\mathbb{P}^{\otimes i}} \{ [F]_n x_{i|k}(W_{0|k}, \dots, \\ & \quad W_{i-1|k}) \} \leq [f]_n, \\ & \forall \quad i \in \mathbb{N}_1^N, n \in \mathbb{N}_1^{n_F}, k \in \mathbb{N}_{\geq 0}. \end{aligned} \quad (5)$$

where $W_{m|k} \sim \mathbb{P}_m$ is the disturbance variable.

3. Distributionally robust optimization and Wasserstein Ambiguity Sets

Distributionally robust optimization is an optimization framework that utilizes partial information about the underlying probability distribution of the random variables in a stochastic model. We consider the Wasserstein ambiguity set (Zhao & Guan, 2018; Mohajerin Esfahani & Kuhn, 2018) in this paper, which is modelled as a Wasserstein ball centered at a discrete empirical distribution. The Wasserstein metric defines the distance between all probability

distributions \mathbb{Q} supported on the uncertainty set $\mathbb{W}_\xi \in \mathbb{R}^{n_\xi}$ with finite p -moment, i.e. $\int_{\mathbb{W}_\xi} \|\xi\|^p \mathbb{Q}(d\xi) < \infty$.

Definition 3.1 (Wasserstein Metric (Ambrosio et al., 2005)). The Wasserstein metric of order $p \geq 1$ is defined as $d_w : \mathcal{M}(\mathbb{W}_\xi) \times \mathcal{M}(\mathbb{W}_\xi) \rightarrow \mathbb{R}$ for all distribution $\mathbb{Q}_1, \mathbb{Q}_2 \in \mathcal{M}(\mathbb{W}_\xi)$ and arbitrary norm on \mathbb{R}^{n_ξ} :

$$d_w^p(\mathbb{Q}_1, \mathbb{Q}_2) := \inf_{\Pi} \int_{\mathbb{W}_\xi^2} \|\xi_1 - \xi_2\|^p \Pi(d\xi_1, d\xi_2), \quad (6)$$

where Π is a joint distribution of ξ_1 and ξ_2 with marginals \mathbb{Q}_1 and \mathbb{Q}_2 respectively.

The Wasserstein metric originates from the optimal transportation problem (Villani, 2009), which studies the most efficient way to transport the mass of a distribution to another. In (6), the Wasserstein distance between the distribution \mathbb{Q}_1 and \mathbb{Q}_2 can be interpreted as the minimal cost spent on the allocation if the Euclidean norm is selected and $p = 2$. In the following, we will regard one distribution as the empirical distribution and the other as one of the unknown distributions which we assess whether to include or not in the ambiguity set. All these unknown distributions, whose distance from the empirical distribution is lower than a certain value in the Wasserstein sense, are included in the ambiguity set.

Specifically, we will only consider the type-1 Wasserstein metric in the remainder of this paper, i.e. $p = 1$. In principle, it is also possible to use other p values given the corresponding reformulation techniques (Kuhn et al., 2019). Then we could define the ambiguity set \mathcal{P} centered at the empirical distribution leveraging the Wasserstein metric as

$$\mathbb{B}_\varepsilon(\hat{\mathbb{P}}) := \left\{ \mathbb{Q} \in \mathcal{M}(\mathbb{W}_\xi) : d_w(\hat{\mathbb{P}}, \mathbb{Q}) \leq \varepsilon \right\} \quad (7)$$

which specifies the Wasserstein ball with radius $\varepsilon > 0$ around the discrete empirical probability distribution $\hat{\mathbb{P}}$. $\mathcal{M}(\mathbb{W}_\xi)$ denotes the set of Borel probability measures on \mathbb{W}_ξ . The empirical probability distribution $\hat{\mathbb{P}} := \frac{1}{M} \sum_{l=1}^M \delta_{\hat{\xi}^{(l)}}$ is the mean of M Dirac distributions which concentrates mass at the disturbance realization $\hat{\xi}^{(l)} \in \mathbb{W}_\xi$. We denote the training set of offline collected realizations ξ as $\hat{\Xi}_M := \left\{ \hat{\xi}^{(l)} \right\}_{l \in \mathbb{N}_1^M} \subset \mathbb{W}_\xi$, which contains M observed disturbance realizations.

The radius ε determines the size of the Wasserstein ball (7), of which the size has been argued from various statistical perspective in the literature (Zhao & Guan, 2018; Rahimian & Mehrotra, 2022; Blanchet et al., 2019). Furthermore, as a function of the radius, the solution of this Wasserstein ambiguity-based DRO lies between the classical robust optimization and sample average approximation (Mohajerin Esfahani & Kuhn, 2018).

4. Propagation of ambiguity sets for nonlinear systems: Iterative distributionally robust LQR

In this section, we propose an algorithm to solve the optimal control problem (5) leveraging the techniques of Wasserstein ambiguity set propagation with dynamic feedback gains and iterative Linear Quadratic Regulator (iLQR) (Li & Todorov, 2004).

For any given sampling time k , we solve the optimal control problem (5) via sequentially deriving the linearized system and the corresponding feedback gains via Riccati recursion (Rawlings et al., 2017, Sec. 8.8.3), solving a perturbed nominal problem, and updating the perturbations based on the solution from a DRO problem. Such an iterative method could be similarly found in, for example, (Messerer & Diehl, 2021). However, our method significantly differs from the robust MPC approach (Messerer & Diehl, 2021) mainly in two aspects: (1) Instead of propagating the state uncertainty based on the ellipsoid support set of additive disturbances, we propagate the Wasserstein ambiguity through the dynamics. (2) We consider soft constraint satisfaction in expectation instead of robust constraint satisfaction.

4.1. LTV formulation and error dynamics

We first consider predicting the system dynamics with the help of tube-based MPC and a linear time-varying (LTV) error system as used in tube-based RMPC (Leeman et al., 2023).

We consider the predicted nonlinear dynamics (1) in the form of the first-order Taylor series expansion:

$$\begin{aligned} x_{i+1|k} &= f_d(x_{i|k}, u_{i|k}) + W_{i|k} \\ &= f_d(z_{i|k}, v_{i|k}) + A(z_{i|k}, v_{i|k})(x_{i|k} - z_{i|k}) \\ &\quad + B(z_{i|k}, v_{i|k})(u_{i|k} - v_{i|k}) \\ &\quad + r(x_{i|k}, u_{i|k}, z_{i|k}, v_{i|k}) + W_{i|k}, \end{aligned} \quad (8)$$

where $A(z_{i|k}, v_{i|k}) := \frac{\partial f_d}{\partial x} \Big|_{(x,u)=(z_{i|k}, v_{i|k})}$, $B(z_{i|k}, v_{i|k}) := \frac{\partial f_d}{\partial u} \Big|_{(x,u)=(z_{i|k}, v_{i|k})}$ and the remainder $r : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \mapsto \mathbb{R}^{n_x}$.

Let $\Delta x_{i|k} := x_{i|k} - z_{i|k}$, $\Delta u_{i|k} := u_{i|k} - v_{i|k}$ denote the errors between nominal and real quantities, we have the following LTV error system

$$\Delta x_{i+1|k} = A_{i|k} \Delta x_{i|k} + B_{i|k} \Delta u_{i|k} + W_{i|k} + r_{i|k}, \quad \Delta x_0 = 0_{n_x}, \quad (9)$$

with $A_{i|k} := A(z_{i|k}, v_{i|k})$, $B_{i|k} := B(z_{i|k}, v_{i|k})$, $r_{i|k} := r(x_{i|k}, u_{i|k}, z_{i|k}, v_{i|k})$. Furthermore, we apply the following control policy with dynamic feedback gain $K_{i|k}$ at each step

$$u_{i|k} = K_{i|k} x_{i|k} + c_{i|k}, \quad (10)$$

where $c_{i|k} \in \mathbb{R}^{n_u}$ are decision variables in the optimal control problem. Also, we have the corresponding nominal policy

$$v_{i|k} = K_{i|k} z_{i|k} + c_{i|k}. \quad (11)$$

Given the error dynamics (9), control policy (10) and nominal policy (11), we have

$$\begin{aligned} \Delta x_{i+2|k} &= A_{cl,i+1} \Delta x_{i+1|k} + W_{i+1|k} + r_{i+1|k} \\ &= A_{cl,i+1} (A_{cl,i} \Delta x_{i|k} + W_{i|k} + r_{i|k}) \\ &\quad + w_{i+1|k} + r_{i+1|k}, \end{aligned} \quad (12)$$

where $A_{cl,i} := A_{i|k} + B_{i|k} K_{i|k}$. Next, let

$$\begin{aligned} e_{i+1|k} &= A_{cl,i} e_{i|k} + W_{i|k} & e_{0|k} &= 0, \\ \varepsilon_{i+1|k} &= A_{cl,i} \varepsilon_{i|k} + r_{i|k} & \varepsilon_{0|k} &= 0. \end{aligned}$$

If $x_{0|k} = z_{0|k}$, i.e. $\Delta x_{0|k} = 0$, we have

$$\Delta x_{i|k} = e_{i|k} + \varepsilon_{i|k}, \forall i \in \mathbb{N}_0^N \quad (13)$$

by induction.

Remark 4.1. In this paper, we consider three cases of feedback gains: fixed feedback with zero gain (open-loop control), fixed feedback with stabilizing gain, and feedback gain computed using iLQR. We will show that if the feedback gain is zero (i.e. open-loop control), the closed-loop performance is significantly worse than with fixed feedback gain and dynamic feedback gain as the size of the propagated ambiguity sets cannot be effectively controlled under open-loop control.

4.2. Distributionally robust nonlinear model predictive control

Given the error dynamics (13), we consider solving an approximated version of the prototype DRNMPC problem (5). In this subsection, similar to (Messerer & Diehl, 2021), we consider the approximated dynamics till the first-order approximation - i.e. ignore the term of linearization error r in (8) - via

$$\begin{aligned} x_{i+1|k} &\approx f_d(z_{i|k}, v_{i|k}) + A_{i|k} (x_{i|k} - z_{i|k}) \\ &\quad + B_{i|k} (u_{i|k} - v_{i|k}) + W_{i|k}. \end{aligned} \quad (14)$$

Based on the approximated linearized dynamics (14), we find the following LTV error dynamics

$$\Delta x_{i+1|k} \approx e_{i+1|k}. \quad (15)$$

Hence the predicted state can be formulated as

$$x_{i+1|k} \approx z_{i+1|k} + A_{cl,i} e_{i|k} + W_{i|k}. \quad (16)$$

Each of the terms evolves under the feedback control (10)

$$\begin{aligned} z_{i+1|k} &= f_d(z_{i|k}, v_{i|k}) & z_{0|k} &= x_k \\ e_{i+1|k} &= A_{cl,i} e_{i|k} + W_{i|k} & e_{0|k} &= 0 \\ v_{i|k} &= K_{i|k} z_{i|k} + c_{i|k}, \end{aligned}$$

where $v_{i|k}$ is the predicted nominal input. As systems' behavior is predicted within a finite prediction horizon, we let $\mathbf{c}_k = [c_{0|k}^\top, \dots, c_{N-1|k}^\top]^\top$ and set $c_{i|k} = 0$ for all $i \geq N$ to ensure a finite number of decision variables. Similarly, we denote \mathbf{v}_k and \mathbf{z}_k for the prediction problem with the horizon N .

Remark 4.2. We will also introduce the error propagation with an explicit consideration of the linearization errors in Section A. However, we will consider only the dynamic ambiguity propagation in this paper for the interest of practical application.

Now we define the objective function for the linearized dynamics at time k as $\mathbb{E}_{\mathbb{P}}\{\sum_{i=0}^{N-1} (\|z_{i|k} + e_{i|k}\|_Q^2 + \|c_{i|k} + K(z_{i|k} + e_{i|k})\|_R^2) + \|z_{N|k} + e_{N|k}\|_{Q_f}^2\}$. Here $Q, Q_f \in \mathbb{R}^{n_x \times n_x}$ and $R \in \mathbb{R}^{n_u \times n_u}$ are positive definite penalty matrices for the quadratic stage costs. Furthermore, based on the assumption of zero-mean additive disturbances, all the accumulated errors $e_{i|k} \forall i \in \mathbb{N}_0^N$ are also zero mean. We could hence reformulate the objective function equivalently to $\sum_{i=1}^{N-1} \|z_{i|k}\|_Q^2 + \|v_{i|k}\|_R^2 + \|z_{N|k}\|_{Q_f}^2$.

Next, we consider the closed-loop propagation of additive disturbances under closed-loop matrices $A_{cl,i}, \forall \mathbb{N}_1^N$. For expected constraints satisfaction (3), we roll out the predicted state ($i \geq 1$) in terms of additive disturbances within the prediction horizon as

$$\begin{aligned} x_{i|k} &\approx z_{i|k} + e_{i|k} \\ &= z_{i|k} + \sum_{m=0}^{i-1} \prod_{j=0}^{i-1-m} A_{cl,i-1-j}^{\min\{1, i-1-m-j\}} W_{m|k}, \end{aligned} \quad (17)$$

where $\prod_{j=0}^{i-1-m} A_{cl,i-1-j}^{\min\{1, i-1-m-j\}} = A_{cl,i-1} \dots A_{cl,i-1-m} I$ for $m < i-1$. Take $x_{3|k}$ as an example, it can be formulated in terms of additive disturbances via $x_{3|k} = A_{cl,2} A_{cl,1} A_{cl,0} x_{0|k} + A_{cl,2} A_{cl,1} w_{0|k} + A_{cl,2} w_{1|k} + w_{2|k}$.

Remark 4.3. The relation

$$e_{i|k} = \sum_{m=0}^{i-1} \prod_{j=0}^{i-1-m} A_{cl,i-1-j}^{\min\{1, i-1-m-j\}} W_{m|k} \quad (18)$$

is the key to our ambiguity set propagation through the (non-linear) dynamics. The most significant difference between our work and existing DRMPC works such as (Mark & Liu, 2020) is that we do not assume having data samples of the predicted states $\hat{x}_{i|k}$, which would simply reduce the optimal control problem to static Wasserstein DRO. However, in practice, one is often faced with the question of having to predict future state distributions and the corresponding ambiguity. We shall demonstrate that, in such dynamic settings, the real power of feedback control is to control the size of the dynamic Wasserstein ambiguity sets in a closed-loop fashion, see Fig. 1. The only previous work considering the

setting equivalent to our dynamic ambiguity set propagation is (Yang, 2020; Aolaritei et al., 2023), but only in the much simpler setting of linear systems. Also, in order to solve dynamic ambiguity set propagation, (Yang, 2020) solves a relaxed problem called the Wasserstein penalty problem (without state constraints), which still requires solving a semi-infinite problem. The technical difficulty that prevents previous works to go beyond that simple setting lies in the very core of Wasserstein DRO reformulation techniques – it does not treat complex nonlinear objectives as in nonlinear OCPs. In contrast, this work proposes the first dynamic Wasserstein closed-loop DRC with nonlinear dynamics and constraints. The Wasserstein distributional reachable set under dynamic propagation will be analytically characterized in Proposition 4.8.

Hence, based on the linearized dynamics, we consider an approximated optimal control problem corresponding to the prototype DRNMPC problem (5).

Problem 2.

$$\begin{aligned}
 \min_{\mathbf{z}, \mathbf{v}, \mathbf{K}} \quad & \sum_{i=0}^{N-1} (\|z_{i|k}\|_Q^2 + \|v_{i|k}\|_R^2) + \|z_{N|k}\|_{Q_f}^2 \\
 \text{s.t.} \quad & z_{0|k} = x_k, \quad z_{i+1|k} = f_d(z_{i|k}, v_{i|k}) \\
 & v_{i|k} = K_{i|k} z_{i|k} + c_{i|k} \\
 & \sup_{\mathbb{P}_m \in \mathcal{P}_{k+m}, m=0, \dots, i-1} \mathbb{E}_{\mathbb{P}^{\otimes i}} \{ [F]_n(z_{i|k} \\
 & \quad + \sum_{m=0}^{i-1} A_{cl}^{(i,j,m)} W_{m|k}) \} \leq [f]_n \\
 \forall \quad & i \in \mathbb{N}_1^N, n \in \mathbb{N}_1^{n_F}, k \in \mathbb{N}_{\geq 0},
 \end{aligned} \tag{19}$$

where $W_{m|k} \sim \mathbb{P}_m$ is the disturbance variable and $A_{cl}^{(i,j,m)} := \prod_{j=0}^{i-1-m} A_{cl, i-1-j}^{\min\{1, i-1-m-j\}}$.

In the following, we will provide the exact reformulation of the optimization problem (19). Before showing the final reformulation, we require the following Lemma to reformulate the distributionally robust constraints.

Lemma 4.4. *Consider the polytopic uncertainty set \mathbb{W}_w and the Wasserstein ambiguity set \mathcal{P} as the Wasserstein ball around the empirical distribution $\hat{\mathbb{P}} = \frac{1}{M} \sum_{l=1}^M \delta_{\hat{w}^{(l)}}$ with type-1 Wasserstein metric and ball radius ε . Then, the worst-case expectation $\sup_{\mathbb{P}_m \in \mathcal{P}} \mathbb{E}_{\mathbb{P}_m} \{ [F]_n(A_{cl}^{(i,j,m)} W_{m|k}) \}$ evaluates to*

$$\begin{aligned}
 \inf_{\lambda, s_l, \gamma_l} \quad & \lambda \varepsilon + \frac{1}{N} \sum_{l=1}^M s_l \\
 \text{s.t.} \quad & [F]_n(A_{cl}^{(i,j,m)} \hat{w}^{(l)}) + \gamma_l^\top (h_w - H_w \hat{w}^{(l)}) \leq s_l \\
 & \|H_w^\top \gamma_l - [F]_n A_{cl}^{(i,j,m)}\|_* \leq \lambda \\
 & \gamma_l \geq 0, \quad \forall l \in \mathbb{N}_1^M,
 \end{aligned} \tag{20}$$

where $\lambda \in \mathbb{R}$, $s_l \in \mathbb{R}$, $\gamma_l \in \mathbb{R}^{n_H}$, and $\|\cdot\|_*$ is the dual norm corresponding to the norm applied in (6).

Proposition 4.5. *Consider the polytopic uncertainty set \mathbb{W}_w . Then, the DRMPC problem (19) evaluate to*

$$\begin{aligned}
 \min_{\lambda_m, s_{ml}, \gamma_{ml}} \quad & \sum_{i=0}^{N-1} (\|z_{i|k}\|_Q^2 + \|v_{i|k}\|_R^2) + \|z_{N|k}\|_{Q_f}^2 \\
 \text{s.t.} \quad & z_{0|k} = x_k \\
 & z_{i+1|k} = f_d(z_{i|k}, v_{i|k}) \\
 & v_{i|k} = K_{i|k} z_{i|k} + c_{i|k} \\
 & \sum_{m=0}^{i-1} \lambda \varepsilon + \frac{1}{N} \sum_{l=1}^M s_{ml} \leq [f]_n - [F]_n z_{i|k} \\
 & [F]_n A_{cl}^{(i,j,m)} \hat{w}^{(l)} + \gamma_l^\top (h_w - H_w \hat{w}^{(l)}) \leq s_{ml} \\
 & \|H_w^\top \gamma_l - [F]_n A_{cl}^{(i,j,m)}\|_* \leq \lambda_m \\
 & \gamma_{ml} \geq 0, \quad \forall l \in \mathbb{N}_1^M, \forall m \in \mathbb{N}_0^{i-1} \\
 & \forall i \in \mathbb{N}_1^N, n \in \mathbb{N}_1^{n_F}, k \in \mathbb{N}_{\geq 0}.
 \end{aligned} \tag{21}$$

Remark 4.6. All the results about the worst-case expected constraints in this paper can be easily replaced by the worst-case chance constraints via the CVaR formulation introduced by (Hota et al., 2019; Xie, 2021). Also, this work can be easily extended to the control problem with polytopic input constraints, e.g. (Zhong et al., 2023).

4.3. Iterative distributionally robust LQR

Optimization problem (19) is difficult to solve as the matrices of the linearized system depend on the unknown nominal trajectory and the unknown system matrices. Hence, we propose in the following an algorithm iterating by sequentially deriving the linearized system and the corresponding feedback gains via Riccati recursion (Rawlings et al., 2017, Sec. 8.8.3), solving a perturbed nominal problem, and updating the perturbations based on the solution from a DRO problem.

We consider below the setting of distributionally robust optimal control, i.e. the problem (21) with the initial sampling time $k = 0$ and a fixed prediction horizon N . Given initial trajectories of nominal state and input \mathbf{z}, \mathbf{v} (e.g. nominal nonlinear MPC), we solve an iterative LQR problem by Riccati Recursion (Mayne, 1966; Li & Todorov, 2004; Rawlings et al., 2017) to get the matrices $\mathbf{A}, \mathbf{B}, \mathbf{K}$ corresponding to the linearized system, where $\mathbf{A} = \{A_{0|k}, \dots, A_{N|k}\}$, $\mathbf{B} = \{B_{0|k}, \dots, B_{N|k}\}$, and $\mathbf{K} = \{K_{0|k}, \dots, K_{N|k}\}$. The feedback gain derived from Riccati recursion makes the closed-loop system matrix A_{cl} stable locally around the linearization point. Then, we solve (20) to update the back-off $\beta = \{\beta_{0|k}, \dots, \beta_{N|k}\}$, where

$$\beta_{i|k} := \sup_{\mathbb{P}_m \in \mathcal{P}, m=0, \dots, i-1} \mathbb{E}_{\mathbb{P}^{\otimes i}} \{ [F]_n (\sum_{m=0}^{i-1} A_{cl}^{(i,j,m)} W_{m|k}) \}. \tag{22}$$

After determining the back-off, we solve the constraint-

tightened program

$$\begin{aligned}
 \min_{\mathbf{z}, \mathbf{v}} \quad & \sum_{i=0}^{N-1} (\|z_{i|k}\|_Q^2 + \|v_{i|k}\|_R^2) + \|z_{N|k}\|_{Q_f}^2 \\
 \text{s.t.} \quad & z_{0|k} = x_k \\
 & z_{i+1|k} = f_d(z_{i|k}, v_{i|k}) \\
 & v_{i|k} = K_{i|k} z_{i|k} + c_{i|k} \\
 & [F]_n z_{i|k} \leq [f]_n - \beta_{i|k}, \\
 & \forall i \in \mathbb{N}_1^N, n \in \mathbb{N}_1^{n_F}, k = 0.
 \end{aligned} \tag{23}$$

We iterate the process above until convergence. The proposed algorithm is summarized in Algorithm 1.

Algorithm 1 Iterative distributionally robust LQR

- 1: **Input:** Initial guess $\mathbf{z}, \mathbf{v}, \varepsilon, \{\hat{w}^{(l)}\}_{l=1}^M$
 - 2: **while** Not Converge **do**
 - 3: $\mathbf{A}, \mathbf{B}, \mathbf{K} \leftarrow$ Riccati Recursion (\mathbf{z}, \mathbf{v}) (Rawlings et al., 2017, Sec. 8.8.3)
 - 4: $\beta \leftarrow$ DRO ($\mathbf{A}, \mathbf{B}, \mathbf{K}, \varepsilon, \{\hat{w}^{(l)}\}_{l=1}^M$) (20)
 - 5: $\mathbf{z}, \mathbf{v}, \mathbf{c} \leftarrow$ Nonlinear optimal control problem with fixed $\mathbf{A}, \mathbf{B}, \mathbf{K}, \beta$ (23)
 - 6: **end while**
 - 7: **Return:** $\mathbf{z}, \mathbf{v}, \mathbf{c}, \mathbf{K}$
-

Remark 4.7. For the original distributionally robust MPC problem (5), we could recursively solve Algorithm 1 at each sampling time k for the given measurement x_k . The initial guess could be generated from a nominal NMPC or a shifted trajectory derived from the previous step.

4.4. Wasserstein distributional reachable sets for error dynamics

We now further show that the stochastic error characterized by the LTV dynamics

$$e_{i+1} = A_{cl,i} e_i + W_i, W_i \stackrel{\text{i.i.d.}}{\sim} \mathbb{P}_w \tag{24}$$

can be contained in a dynamic Wasserstein ambiguity set given below. This gives a theoretical bound for the experimental results in Figure 1.

Let $\hat{e}_i \sim \hat{Q}_{N,i}$ be the empirical error vector at predicted time step i , and $\hat{Q}_{N,i}$ be its empirical error distribution, given by the empirical LTV error dynamics

$$\hat{e}_{i+1} = A_{cl,i} \hat{e}_i + \hat{W}_i, e_0 = 0, \hat{W}_i \stackrel{\text{i.i.d.}}{\sim} \hat{\mathbb{P}}.$$

We now consider a dynamic ambiguity set—*Wasserstein ambiguity tube* associated with the LTV dynamics, with a slight abuse of notation,

$$\mathbb{T}_\epsilon(\hat{P}_N) := \left\{ Q_i, i = 0, \dots, T \mid e_i \sim Q_i, e_{i+1} = A_{cl,i} e_i + W_i, W_i \sim \mathbb{P}_w, d_w^p(\mathbb{P}_w, \hat{\mathbb{P}}) \leq \epsilon, e_0 = 0 \right\}.$$

The intuition is that the Wasserstein ambiguity tube contains all evolution paths of the ambiguous stochastic system, i.e., the state distribution of our MPC problem lives in this ambiguity tube $\{Q_i\} \in \mathbb{T}_\epsilon(\hat{\mathbb{P}})$.

The following result characterizes the size of the Wasserstein ambiguity tube.

Proposition 4.8 (Wasserstein distributional reachable sets). *We have, $\forall \{Q_i\} \in \mathbb{T}_\epsilon(\hat{\mathbb{P}})$,*

$$d_w^p(Q_i, \hat{Q}_{N,i}) \leq \epsilon \cdot \sum_{m=0}^{i-1} \prod_{j=0}^{i-1-m} \|A_{cl,i-1-j}^{\min\{1, i-1-m-j\}}\|^p,$$

for $i = 0, \dots, T$, where $\|\cdot\|$ is the corresponding induced matrix norm.

Proposition 4.8 equips us with a reachable set for the error in the Wasserstein distance. We illustrate this in a Figure 1. It further implies that the aforementioned ambiguity tube can be bounded in a more straightforward and computable dynamic ambiguity set (tube)

$$\begin{aligned}
 \mathbb{T}_\epsilon(\hat{\mathbb{P}}) \subset \left\{ Q_i, i = 0, \dots, T \mid d_w^p(Q_i, \hat{Q}_{N,i}) \leq \right. \\
 \left. \epsilon \cdot \sum_{m=0}^{i-1} \prod_{j=0}^{i-1-m} \|A_{cl,i-1-j}^{\min\{1, i-1-m-j\}}\|^p, \hat{e}_i \sim \hat{Q}_{N,i}, \right. \\
 \left. \hat{e}_{i+1} = A_{cl,i} \hat{e}_i + \hat{W}_i, \hat{W}_i \stackrel{\text{i.i.d.}}{\sim} \hat{\mathbb{P}}, e_0 = 0. \right\} \tag{25}
 \end{aligned}$$

Unlike aforementioned works in the existing literature where the ambiguity sets are often given a priori, our dynamic Wasserstein ambiguity set is obtained by propagating through the LTV error dynamics.

Remark 4.9 (Wasserstein invariant ambiguity sets). While this paper does not deal with infinite-horizon control or positive invariant sets, it is easy to see that Proposition 4.8 can be used to construct Wasserstein invariant ambiguity sets by examining the series $\sum_{m=0}^{\infty} \prod_{j=0}^m \|A_{cl,i-1-j}^{\min\{1, i-1-m-j\}}\|^p$, i.e., if $\exists C < \infty$ such that $\sum_{m=0}^{\infty} \prod_{j=0}^m \|A_{cl,i-1-j}^{\min\{1, i-1-m-j\}}\|^p \leq C$, then the following set of distributions is a Wasserstein invariant set for the ambiguous system state distribution $\left\{ Q \mid d_w^p(Q, \hat{Q}_{N,\infty}) \leq \epsilon \cdot C \right\}$, where $\hat{Q}_{N,\infty}$ is the equilibrium state distribution of the nominal error dynamics (24).

Remark 4.10. The error dynamics propagation for linearization errors can be found in Section A.

5. Case study

The system considered is a nonlinear mass spring system with $m = 2$ kg, $k_1 = 3$ N/m, $k_2 = 2$ N/m:

$$\dot{x}_1 = x_2, \quad \dot{x}_2 = -\frac{k_2}{m} x_1^5 - \frac{k_1}{m} x_2 + \frac{1}{m} u.$$

The discrete-time system is acquired by using the Runge-Kutta method with fourth order with the sampling period 0.1 s. We simulate the control performance for the discrete-time system suffering from the uniformly distributed additive disturbance bounded within $[-1e-3, 1e-3]$ on the state element x_1 , and $[-0.1, 0.1]$ on x_2 . The prediction horizon for this system is set to $N = 140$.

The control goal of this system is to track the state $x_r = [0, 0]^T$ starting from the initial state $x_{init} = [-2, 0]^T$, while satisfying the distributionally robust state constraint corresponding to $x_2 \leq 0.5$ m/s. The parameters are selected as

$$Q = Q_f = \begin{bmatrix} 100 & 0 \\ 0 & 1 \end{bmatrix}, R = [1].$$

We compare the closed-loop performance of three different methods: our method, fixed feedback gain, and no feedback gain in this section. The fixed feedback gain $K = [-7.97, -7.16]$ is derived from the LQR controller for the nonlinear system linearized around the equilibrium point with the same penalty matrices.

With $M = 5$ offline collected disturbance samples and ball radius $\varepsilon = 0.03$, simulation results of nominal trajectories solved by Algorithm 1 with the three different methods mentioned above (for the method with the fix or zero feedback gain we only linearize the nominal nonlinear system without update K) can be found in Fig. 2. We apply $u_k = K_k x_k + c_k$ to the disturbed nonlinear system from $k = 0$ to $k = N$, where $K_k = 0$ for the method with zero feedback gain and $K_k = [-7.97, -7.16]$ for the method with fixed feedback gain. Fig 3 illustrates 20 realizations of the closed-loop performance. Fig. 1 illustrates the accumulated error between the closed-loop state and the nominal state shown in Fig. 2. The arrow indicates the error difference between two consecutive sampling times, i.e. the tail indicates the accumulated error with respect to the nominal state of the previous steps and the head indicates the accumulated error with respect to the nominal state at the current step.

Based on our experiments, we conclude that though the open-loop controller is capable of controlling the nominal trajectory, it is not effective in reducing the growth of ambiguity, as depicted in Fig 1. Conversely, our feedback controller successfully controls the size of ambiguity, which is the main insight of our paper.

6. Conclusions

Our key insight is that the sizes of Wasserstein ambiguity sets for nonlinear systems can be controlled using nonlinear feedback control. To demonstrate that, this paper proposes the DRNMPC with dynamic Wasserstein ambiguity. We propose an iterative MPC scheme to dynamically control the propagation of Wasserstein ambiguity sets. We analytically characterize the Wasserstein distributional reachable set un-

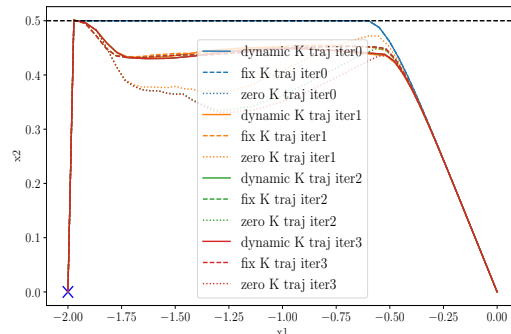


Figure 2. Nominal state trajectories solved by Algorithm 1 with Wasserstein ball radius $\varepsilon = 0.03$.




Figure 3. Closed-loop performance of 30 realizations with the feedback gains and nominal inputs solved by 1. Red: Out method. Blue: Fixed feedback gain. Black: Zero feedback gain.

der dynamic propagation in our algorithm. To evaluate the effectiveness of our proposed algorithm, we compare the closed-loop performances of dynamic feedback, fixed feedback, and no feedback on a mass-spring system. The simulation results demonstrate that the proposed iterative scheme can effectively control the ambiguity set propagation, which is a critical step in solving the DRNMPC problem.

References

- Ambrosio, L., Gigli, N., and Savaré, G. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2005.
- Aolaritei, L., Lanzetti, N., and Dörfler, F. Capture, propagate, and control distributional uncertainty. *arXiv preprint arXiv:2304.02235*, 2023.
- Blanchet, J., Kang, Y., and Murthy, K. Robust wasserstein profile inference and applications to machine learning. *Journal of Applied Probability*, 56(3):830–857, 2019.
- Cannon, M., Kouvaritakis, B., and Wu, X. Model predictive control for systems with stochastic multiplicative uncer-

- tainty and probabilistic constraints. *Automatica*, 45(1):167–172, 2009.
- Coulson, J., Lygeros, J., and Dorfler, F. Distributionally robust chance constrained data-enabled predictive control. *IEEE Transactions on Automatic Control*, 2021.
- Fochesato, M. and Lygeros, J. Data-driven distributionally robust bounds for stochastic model predictive control. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 3611–3616. IEEE, 2022.
- Hakobyan, A. and Yang, I. Wasserstein distributionally robust control of partially observable linear stochastic systems. *arXiv preprint arXiv:2212.04644*, 2022.
- Heirung, T. A. N., Paulson, J. A., O’Leary, J., and Mesbah, A. Stochastic model predictive control—how does it work? *Computers & Chemical Engineering*, 114:158–170, 2018.
- Hota, A. R., Cherukuri, A., and Lygeros, J. Data-driven chance constrained optimization under wasserstein ambiguity sets. In *2019 American Control Conference (ACC)*, pp. 1501–1506. IEEE, 2019.
- Kuhn, D., Esfahani, P. M., Nguyen, V. A., and Shafieezadeh-Abadeh, S. Wasserstein distributionally robust optimization: Theory and applications in machine learning. In *Operations research & management science in the age of analytics*, pp. 130–166. Informa, 2019.
- Leeman, A. P., Köhler, J., Zanelli, A., Bennani, S., and Zeilinger, M. N. Robust nonlinear optimal control via system level synthesis. *arXiv preprint arXiv:2301.04943*, 2023.
- Li, W. and Todorov, E. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *ICINCO (1)*, pp. 222–229. Citeseer, 2004.
- Mark, C. and Liu, S. Stochastic mpc with distributionally robust chance constraints. *IFAC-PapersOnLine*, 53(2):7136–7141, 2020.
- Mayne, D. A second-order gradient method for determining optimal trajectories of non-linear discrete-time systems. *International Journal of Control*, 3(1):85–95, 1966.
- Mayne, D. Q. Model predictive control: Recent developments and future promise. *Automatica*, 50(12):2967–2986, 2014.
- Mayne, D. Q., Rawlings, J. B., Rao, C. V., and Scokaert, P. O. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.
- Mayne, D. Q., Seron, M. M., and Raković, S. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224, 2005.
- Messerer, F. and Diehl, M. An efficient algorithm for tube-based robust nonlinear optimal control with optimal linear feedback. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pp. 6714–6721. IEEE, 2021.
- Micheli, F., Summers, T., and Lygeros, J. Data-driven distributionally robust mpc for systems with uncertain dynamics. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pp. 4788–4793. IEEE, 2022.
- Mohajerin Esfahani, P. and Kuhn, D. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166, 2018.
- Rahimian, H. and Mehrotra, S. Frameworks and results in distributionally robust optimization. *Open Journal of Mathematical Optimization*, 3:1–85, 2022.
- Rawlings, J. B., Mayne, D. Q., and Diehl, M. *Model predictive control: theory, computation, and design*, volume 2. Nob Hill Publishing Madison, WI, 2017.
- Villani, C. *Optimal transport: old and new*, volume 338. Springer, 2009.
- Xie, W. On distributionally robust chance constrained programs with wasserstein distance. *Mathematical Programming*, 186(1-2):115–155, 2021.
- Yang, I. Wasserstein distributionally robust stochastic control: A data-driven approach. *IEEE Transactions on Automatic Control*, 66(8):3863–3870, 2020.
- Zhao, C. and Guan, Y. Data-driven risk-averse stochastic optimization with wasserstein metric. *Operations Research Letters*, 46(2):262–267, 2018.
- Zhong, Z., del Rio-Chanona, E. A., and Petsagkourakis, P. Data-driven distributionally robust mpc using the wasserstein metric. *arXiv preprint arXiv:2105.08414*, 2021.
- Zhong, Z., del Rio-Chanona, E. A., and Petsagkourakis, P. Tube-based distributionally robust model predictive control for nonlinear process systems via linearization. *Computers & Chemical Engineering*, 170:108112, 2023.
- Zolanvari, A. and Cherukuri, A. Data-driven distributionally robust iterative risk-constrained model predictive control. In *2022 European Control Conference (ECC)*, pp. 1578–1583. IEEE, 2022.

A. Linearization error reachable sets

Let us now consider the error dynamics propagation for linearization errors with the following standard assumption.

Assumption A.1. The nonlinear dynamics (1) $f_d : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \mapsto \mathbb{R}^{n_x}$ are three times continuously differentiable.

To bound the linearization error, we consider the following condition of locally bounded eigenvalues on Hessian, similar as in (Leeman et al., 2023). Let $\mathcal{X} \in \mathbb{R}^{n_x}$ and $\mathcal{U} \in \mathbb{R}^{n_u}$ denote the state and input space, respectively and $H_n : \mathbb{R}^{n_x+n_u} \mapsto \mathbb{R}^{(n_x+n_u) \times (n_x+n_u)}$ denote the Hessian corresponding to the n-th component of f_d , i.e.

$$H_n(\xi_x, \xi_u) = \begin{bmatrix} \frac{\partial^2 f_{d,n}}{\partial x^2} & \frac{\partial^2 f_{d,n}}{\partial x \partial u} \\ * & \frac{\partial^2 f_{d,n}}{\partial u^2} \end{bmatrix} \Big|_{(x,u)=(\xi_x, \xi_u)},$$

where $\xi_x \in \mathcal{X}$ and $\xi_u \in \mathcal{U}$. We further denote the constant μ_n as the corresponding locally maximal eigenvalue, i.e.

$$\mu_n := \frac{1}{2} \max_{\xi_x \in \mathcal{X}, \xi_u \in \mathcal{U}, \|h\|_\infty \leq 1} |h^\top H_n(\xi_x, \xi_u) h|.$$

Then we have the following bound for each n-th element of the vector of linearization errors $r(x, u, z, v)$ in (8).

Lemma A.2. (Leeman et al., 2023, Proposition III.1.) *Given Assumption A.1, the remainder in (8) satisfies*

$$|r_n(x, u, z, v)| \leq \|\eta\|_\infty^2 \mu_n,$$

for any $z, x \in \mathcal{X}, u, v \in \mathcal{U}$, where $\eta = \begin{bmatrix} \Delta x \\ \Delta u \end{bmatrix}$.

Proof. By second-order Mean Value Theorem, we know that there exist $\xi_x \in [x, z]$ and $\xi_u \in [u, v]$ (with a little abuse of notation) such that $r_n(x, u, z, v) = \frac{1}{2} \begin{bmatrix} x - z \\ u - v \end{bmatrix}^\top H_n(\xi_x, \xi_u) \begin{bmatrix} x - z \\ u - v \end{bmatrix}$. Hence we have

$$\begin{aligned} |r_n(x, u, z, v)| &\leq \max_{\xi_x, \xi_u} \frac{1}{2} \left| \begin{bmatrix} \Delta x \\ \Delta u \end{bmatrix}^\top H_n(\xi_x, \xi_u) \begin{bmatrix} \Delta x \\ \Delta u \end{bmatrix} \right| \\ &\leq \left\| \begin{bmatrix} \Delta x \\ \Delta u \end{bmatrix} \right\|_\infty^2 \max_{\xi_x \in \mathcal{X}, \xi_u \in \mathcal{U}, \|h\| \leq 1} \frac{1}{2} |h^\top H_n(\xi_x, \xi_u) h|. \end{aligned}$$

If the infinity norm is considered here, we have $|r_n(x, u, z, v)| \leq \|\eta\|_\infty^2 \mu_n$. \square

Given Lemma A.2, dynamics (13) and the control policy, we have the following lemma characterizing the upper bound of the linearization error.

Lemma A.3. *Given Assumption A.1, the remainder in (8) is upper bounded by*

$$r_n(x_{k+i}, u_{k+i}, z_{k+i}, v_{k+i}) \leq \left(\left\| \begin{bmatrix} e_{k+i} \\ K_{k+i} e_{k+i} \end{bmatrix} \right\|_\infty + \left\| \begin{bmatrix} \varepsilon_{k+i} \\ K_{k+i} \varepsilon_{k+i} \end{bmatrix} \right\|_\infty \right)^2 \mu_n, \quad (26)$$

for any $z, x \in \mathcal{X}, u, v \in \mathcal{U}$, where $i \in \mathbb{N}_+$ and $n \in \mathbb{N}_1^{n_x}$.

Proof. This can be directly derived from Lemma A.2 with the error dynamics (13). \square

Furthermore, let $\mu := \text{diag}(\mu_1, \dots, \mu_{n_x})$, then the linearization error satisfies

$$r_{k+i} \in \left(\left\| \begin{bmatrix} e_{k+i} \\ K_{k+i} e_{k+i} \end{bmatrix} \right\|_\infty + \left\| \begin{bmatrix} \varepsilon_{k+i} \\ K_{k+i} \varepsilon_{k+i} \end{bmatrix} \right\|_\infty \right)^2 \mu \mathbb{B}_\infty^{n_x},$$

where $r_{k+i} := \begin{bmatrix} r_1(x_{k+i}, u_{k+i}, z_{k+i}, v_{k+i}) \\ \vdots \\ r_{n_x}(x_{k+i}, u_{k+i}, z_{k+i}, v_{k+i}) \end{bmatrix}$.

Remark A.4. The back-off due to the linearization error can be derived from Lemma A.3. However, we observed that, in practical experiments, the linearization error propagation might result in an over-conservative closed-loop performance; hence, we will only consider the dynamic propagation of ambiguity sets in practical numerical experiments below.

B. Proofs

B.1. Proof to Lemma 4.4

Proof. The equivalent reformulation can be derived with $a_k := [F]_n A_{cl}^{(i,j,m)}$ in (Mohajerin Esfahani & Kuhn, 2018, Corollary 5.1). \square

B.2. Proof to Proposition 4.5

Proof. The reformulation (21) is the consequence of the exact reformulation of the distributionally robust constraints. For any given i and k , based on the linearity property of expectation, the distributionally robust constraints $\sup_{r_{m|k}, \mathbb{P}_0 \in \mathcal{P}, \dots, \mathbb{P}_{i-1} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}^{\otimes i}} \{ [F]_n (z_{i|k} + \sum_{m=0}^{i-1} A_{cl}^{(i,j,m)} W_{m|k}) \} \leq [f]_n$ is equivalent to

$$\sum_{m=0}^{i-1} \sup_{\mathbb{P}_m \in \mathcal{P}} \mathbb{E}_{\mathbb{P}_m} \{ [F]_n (A_{cl}^{(i,j,m)} W_{m|k}) \} + [F]_n z_{i|k} \leq [f]_n. \quad (27)$$

Let $\mathbb{Z}_f := \{ \mathbf{z} \mid \mathbf{z} \text{ is feasible in (27) } \forall i \in \mathbb{N}_1^N \}$. By applying Lemma 4.4 to $\sup_{\mathbb{P}_m \in \mathcal{P}} \mathbb{E}_{\mathbb{P}_m} \{ [F]_n (A_{cl}^{(i,j,m)} W_{m|k}) \}$ for each m , we acquire that the inequality (27) containing the summation of distributionally robust optimizations is equivalent to

$$\begin{aligned} \sum_{m=0}^{i-1} \inf_{\lambda_m, s_{ml}, \gamma_{ml}} \lambda_m \varepsilon + \frac{1}{N} \sum_{l=1}^M s_{ml} &\leq [f]_n - [F]_n z_{i|k} \\ \text{s.t. } [F]_n A_{cl}^{(i,j,m)} \hat{w}_l + \gamma_l^\top (h - H \hat{w}_l) &\leq s_{ml} \\ \|H_w^\top \gamma_l - [[F]_n A_{cl}^{(i,j,m)}]^\top\|_* &\leq \lambda_m \\ \gamma_{ml} &\geq 0, \quad \forall l \in \mathbb{N}_1^M, \forall m \in \mathbb{N}_0^{i-1}. \end{aligned} \quad (28)$$

Hence the feasible set \mathbb{Z}_f is equivalent to

$$\begin{aligned} \mathbb{Z}_f := \{ \mathbf{z} \mid \exists \lambda_m, s_{ml}, \gamma_{ml} \text{ s.t. } \mathbf{z} \text{ is feasible in} \\ \sum_{m=0}^{i-1} \lambda_m \varepsilon + \frac{1}{N} \sum_{l=1}^M s_{ml} &\leq [f]_n - [F]_n z_{i|k} \\ [F]_n A_{cl}^{(i,j,m)} \hat{w}_l + \gamma_l^\top (h - H \hat{w}_l) &\leq s_{ml} \\ \|H_w^\top \gamma_l - [[F]_n A_{cl}^{(i,j,m)}]^\top\|_* &\leq \lambda_m \\ \gamma_{ml} &\geq 0, \quad \forall l \in \mathbb{N}_1^M, \forall m \in \mathbb{N}_0^{i-1} \\ \forall i \in \mathbb{N}_1^N \}. \end{aligned}$$

Together with the objective function and equality constraints corresponding to the nominal dynamic, we complete the proof. \square

B.3. Proof to Proposition 4.8

Proof. By the definition of the Wasserstein distance,

$$d_w^p(Q_i, \hat{Q}_{N,i}) = \inf_{\Pi} \int_{\mathbb{W}_w^2} \|e_i^* - \hat{e}_i\|^p d\Pi(e_i^*, \hat{e}_i),$$

where Π is a joint distribution (transport plan) with marginals $Q_i, \hat{Q}_{N,i}$.

Plugging in the error dynamics decomposition (18),

$$\begin{aligned}
 & \inf_{\Gamma} \int_{\mathbb{W}_w^2} \left\| \sum_{m=0}^{i-1} \prod_{j=0}^{i-1-m} A_{cl, i-1-j}^{\min\{1, i-1-m-j\}} (W_m^* - \hat{W}_m) \right\|^p d\Gamma (W_m^*, \hat{W}_m) \\
 & \leq \sum_{m=0}^{i-1} \left\| \prod_{j=0}^{i-1-m} A_{cl, i-1-j}^{\min\{1, i-1-m-j\}} \right\|^p \cdot \inf_{\Gamma} \int_{\mathbb{W}_w^2} \|W_m^* - \hat{W}_m\|^p d\Gamma (W_m^*, \hat{W}_m) \\
 & \leq \epsilon \cdot \sum_{m=0}^{i-1} \prod_{j=0}^{i-1-m} \left\| A_{cl, i-1-j}^{\min\{1, i-1-m-j\}} \right\|^p,
 \end{aligned}$$

where Γ is the joint distribution of W_m^*, \hat{W}_m . Due to the dynamics structure (24), the joint distribution Π is determined by the joint distributions Γ . The last inequality estimate above is due to the Wasserstein distance estimate $d_w^p(\mathbb{P}_w, \hat{\mathbb{P}}) \leq \epsilon$. \square