

Data-driven distributionally robust MPC for systems with uncertain dynamics

Conference Paper

Author(s): Micheli, Francesco; Summers, Tyler; Lygeros, John

Publication date: 2022

Permanent link: https://doi.org/10.3929/ethz-b-000584148

Rights / license: In Copyright - Non-Commercial Use Permitted

Originally published in: https://doi.org/10.1109/CDC51059.2022.9992469

Funding acknowledgement: 787845 - Optimal control at large (EC)

Data-driven distributionally robust MPC for systems with uncertain dynamics

Francesco Micheli, Tyler Summers, John Lygeros

Abstract—We present a novel data-driven distributionally robust Model Predictive Control formulation for unknown discrete-time linear time-invariant systems affected by unknown and possibly unbounded additive uncertainties. We use off-line collected data and an approximate model of the dynamics to formulate a finite-horizon optimization problem. To account for both the uncertainty related to the dynamics and the disturbance acting on the system, we resort to a distributionally robust formulation that optimizes the cost expectation while satisfying Conditional Value-at-Risk constraints with respect to the worst-case probability distributions of the uncertainties within an ambiguity set defined using the Wasserstein metric. Using results from the distributionally robust optimization literature we derive a tractable finite-dimensional convex optimization problem with finite-sample guarantees for the class of convex piecewise affine cost and constraint functions. The performance of the proposed algorithm is demonstrated in closed-loop simulation on a simple numerical example.

I. INTRODUCTION

Model Predictive Control (MPC) relies on a model of the system dynamics to repeatedly solve, at each time-step, a Finite-Horizon Optimal Control (FHOC) problem subject to input and state constraints. Control performance is related to the quality of the open-loop predictions, making MPC susceptible to uncertainties in the prediction model. Uncertainty is present in many practically relevant control applications, as both the dynamics and the disturbance distribution are only approximately known or must be estimated from data. Hence, in the last two decades, robust and stochastic MPC have been developed to cope explicitly with uncertainty whenever the robustness that might be implicitly provided by feedback from deterministic MPC is insufficient.

Robust MPC (RMPC) [1] relies on the assumption of bounded uncertainties to solve a worst-case optimization problem. Tube-based formulations [2], [3] have been developed to account for both the uncertainty on the model and on the disturbance realizations. Since RMPC accounts for all possible realization of the uncertainties, thus neglecting any available distributional information, it often results in a conservative controller.

When distributional information on the model uncertainty and disturbance is available, Stochastic MPC (SMPC) can reduce the conservatism of RMPC by enforcing constraints to hold in probability [4], [5]. Unfortunately, analytical solutions to SMPC are only available under specific assumptions on the uncertainties distributions. Alternatively, randomized methods such as the sample average approximation [6] and the scenario approach [7] can be employed to reformulate the stochastic problem into a large, but finite dimensional, deterministic program. Since randomized methods leverage sampling, they can handle generic distributions and can be applied when these are only accessible through sampling. Applications of the scenario approach for a SMPC with uncertain dynamics and unbounded stochastic disturbance are analyzed in [8] and [9], where the scenarios are obtained by sampling the dynamics and the disturbance distributions.

Crucially, RMPC and SMPC rely on specific assumptions on the disturbances or, in the case of the randomized methods, availability of large amount of data is required to provide tight probabilistic guarantees. In many situations, we might only have access to an approximate dynamics and limited amount of data regarding disturbance process, resulting in uncertainty on the dynamics and on disturbance distribution. We therefore resort to a Distributionally Robust MPC (DRMPC) formulation that solves a stochastic optimization problem with respect to the worst-case probability distribution within an ambiguity set. This allows to avoid the excessive conservatism of RMPC while protecting against the distributional mismatch that can negatively impact SMPC.

Lately, a number of works developed DRMPC formulation to robustify against additive disturbances or uncertain constraints considering moment-based [10], [11], [12] and Wasserstein-based [13], [14], [15] uncertainty sets. None of these approaches directly addresses the issue of uncertainty in the system dynamics parameters; they instead focus on a limited uncertainty representation from additive noise only. This is a fundamental aspect to consider whenever the system is too complex to be precisely modeled and an approximate model is obtained from limited data or when the controller needs to be robust against plant changes resulting from production variability or system aging.

In this work we propose a data-driven DRMPC formulation that robustifies against both the uncertainty on the approximate dynamics and the uncertainty on the additive disturbance acting on the system. Unlike earlier works on DRMPC, we directly robustify against the distributional mismatch in the predicted state trajectories. We do so by defining the ambiguity set as a ball centered on the empirical distribution given by the predictions obtained with the uncertain dynamics and radius defined with respect to

Research supported by the European Research Council under the H2020 Advanced Grant no. 787845 (OCAL).

F. Micheli and J. Lygeros are with the Automatic Control Laboratory in the Department of Information Technology and Electrical Engineering, ETH Zürich, Switzerland. Emails: {frmicheli, jlygeros}@ethz.ch.

T. Summers is with the Department of Mechanical Engineering at University of Texas at Dallas, USA. Email: tyler.summers@utdallas.edu. The work of T. Summers is supported by the United States Air Force Office of Scientific Research under award number FA2386-19-1-4073 and the National Science Foundation under award number ECCS-2047040.

the Wasserstein metric. This allows us to capture both the uncertainty in the approximate dynamics and the uncertainty on the disturbance distributions. Leveraging earlier results from the distributionally robust literature we derive finite sample guarantees and present tractable finite-dimensional convex formulations for the DRMPC problem with worstcase expectation cost and Conditional Value-at-Risk constraints. The derived formulation generalizes previous works on DRMPC that only address the effect of the additive disturbance and consider known dynamics. Whenever the dynamics can be considered exactly known, the formulation reduces to a DRMPC with only additive uncertainty, highlighting the consistency of the approach.

Outline: The remainder of this paper is organized as follows. In Section II we introduce the problem setting and formally state the distributionally robust control problem. In Section III we derive finite sample guarantees for the DRMPC problem with worst-case expectation cost and Conditional Value-at-Risk constraints. Section IV demonstrates the effectiveness of the proposed algorithm on a simple closed-loop example and Section V concludes the paper.

Notation: We denote by δ_x the Dirac distribution at x. We denote by $\|\cdot\| := \|\cdot\|_2$ the standard Euclidean distance, for matrix norms the same symbol is used to denote the norm induced by the 2-norm. We define by $(\cdot)_+ := \max\{\cdot, 0\}$.

II. PROBLEM FORMULATION

Consider the discrete-time linear time-invariant (LTI) system subject to additive disturbances

$$x_{k+1} = Ax_k + Bu_k + w_k \; ,$$

with state $x_k \in \mathbb{R}^n$, control input $u_k \in \mathbb{R}^m$ and disturbance $w_k \in \mathbb{R}^n$ distributed according to a unknown probability distribution \mathbb{P}_w over the unknown and possibly unbounded support set $\mathcal{W} \subseteq \mathbb{R}^n$. We assume that the state is measured, but that we do not have access to the true system dynamics.

We introduce the compact T-step formulation of the dynamics

$$\boldsymbol{y} = \bar{L}\boldsymbol{z} + \boldsymbol{\xi} \tag{1}$$

with

$$\boldsymbol{y} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_T \end{bmatrix}, \quad \boldsymbol{u} = \begin{bmatrix} u_0 \\ u_1 \\ \vdots \\ u_{T-1} \end{bmatrix}, \quad \boldsymbol{w} = \begin{bmatrix} w_0 \\ w_1 \\ \vdots \\ w_{T-1} \end{bmatrix}, \quad \boldsymbol{z} = \begin{bmatrix} \boldsymbol{x}_0 \\ \boldsymbol{u} \end{bmatrix}, \quad \boldsymbol{\xi} = \bar{H}\boldsymbol{w},$$

$$\bar{L} = \begin{bmatrix} \bar{A} & \bar{B} & 0_{n\times m} & \cdots & 0_{n\times m} \\ \bar{A}^2 & \bar{A}\bar{B} & \bar{B} & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0_{n\times m} \\ \bar{A}^T & \bar{A}^{T-1}\bar{B} & \cdots & \bar{A}\bar{B} & \bar{B} \end{bmatrix}, \quad \bar{H} = \begin{bmatrix} I_{n\times n} & 0_{n\times n} & \cdots & 0_{n\times n} \\ \bar{A} & I_{n\times n} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0_{n\times n} \\ \bar{A}^{T-1} & \cdots & \bar{A} & I_{n\times n} \end{bmatrix}$$

The future state sequence y depends on the initial condition x_0 , the control input sequence u and the multi-step disturbance $\boldsymbol{\xi} \in \Xi \subseteq \mathbb{R}^{nT}$ distributed according to $\mathbb{P}_{\boldsymbol{\xi}}$. We also define $\mathbb{P}_{\boldsymbol{y}|\boldsymbol{z}}$ as the probability distribution of \boldsymbol{y} resulting from the dynamics (1) for given initial condition and control input sequence.

We are interested in designing a receding horizon predictive controller that minimizes a given cost function over a finite horizon T. At each time τ we would like to solve the finite horizon optimal control (FHOC) problem

$$= \min_{\boldsymbol{z}} \quad \mathbb{E}^{\boldsymbol{y} \sim \mathbb{E}^{\boldsymbol{y} \mid \boldsymbol{z}} \mid \boldsymbol{y} \mid \boldsymbol{z}} \left[h\left(\boldsymbol{y}, \boldsymbol{z} \right) \right]$$
s.t. $\boldsymbol{z} \in \mathcal{Z}_{\tau}$

$$\text{CVaR}_{\boldsymbol{1} - \boldsymbol{\beta}}^{\boldsymbol{y} \sim \mathbb{P}^{\boldsymbol{y} \mid \boldsymbol{z}}} \left(g(\boldsymbol{y}, \boldsymbol{z}) \right) \leq 0 .$$

$$(2)$$

The set \mathcal{Z}_{τ} includes a constraint that sets \boldsymbol{x}_0 equal to the measured value of the state at the time τ when (2) is solved, and possibly other constraints that we want to impose on the input trajectory. Since the future states \boldsymbol{y} are uncertain, we enforce the constraint $g(\boldsymbol{y}, \boldsymbol{z}) \leq 0$ with $g(\boldsymbol{y}, \boldsymbol{z}) : \mathbb{R}^{nT} \times \mathbb{R}^{nT} \to \mathbb{R}$ as a Conditional Value-at-Risk (CVaR) constraint.

Definition 1 (Conditional Value-at-Risk). For a random variable $\omega \in \Omega \subseteq \mathbb{R}^r$ with distribution \mathbb{P}_{ω} and a function $\phi : \mathbb{R}^r \to \mathbb{R}$, the CVaR of level β is defined as

$$CVaR_{1-\beta}^{\omega \sim \mathbb{P}_{\omega}}\left(\phi(\omega)\right) := \inf_{t \in \mathbb{R}} \left[\beta^{-1} \mathbb{E}^{\omega \sim \mathbb{P}_{\omega}}\left[\left(\phi(\omega) + t\right)_{+}\right] - t\right] \;.$$

A. The sample average approximation

 \bar{J}_{τ}

Since \bar{L} and $\mathbb{P}_{\boldsymbol{\xi}}$ are unknown, (2) cannot be solved directly. We assume the availability of a dataset $\mathcal{D}^{N,T}$ comprising N T-step input-state trajectories $\{\boldsymbol{z}^i, \boldsymbol{y}^i\}, i = 1, \ldots, N$, collected by applying an input trajectory \boldsymbol{u}^i of length T starting from an initial condition \boldsymbol{x}_0^i . We also assume to have access to an approximate T-step linear predictor \hat{L} , obtained from a separate identification procedure, and make the following assumption

Assumption 2. For a given confidence level $\alpha \in (0, 1)$, let $\gamma(\alpha) : (0, 1) \to \mathbb{R}$ be a function such that

$$\mathbb{P}^{\mathcal{ID}}\left\{\left\|\bar{L}-\hat{L}\right\| \leq \gamma(\alpha)\right\} \geq 1-\alpha \ .$$

With $\mathbb{P}^{\mathcal{ID}}$ the probability related to the identification process.

With this assumption we essentially bound the mismatch between the approximate and true dynamics. References [16], [17], [18] provide conditions for Assumption 2 to hold, when the prediction model is obtained by linear regression from historical data. Nonetheless, the resulting theoretical bounds can be quite conservative in practice [16].

We can use the data in $\mathcal{D}^{N,T}$ and \hat{L} to compute N approximate T-step residuals

$$\hat{\xi}^{i} := y^{i} - \hat{L}z^{i} , \quad i = 1, \dots, N .$$
 (3)

Given some initial condition x_0 and control input trajectory u, we can use the multi-step predictor \hat{L} and the multi-step residuals obtained in (3) to compute approximate multi-step predictions as

$$\hat{y}^{i}(z) := \hat{L}z + \hat{\xi}^{i}, \quad i = 1, \dots, N$$
 (4)

In the spirit of the Sample Average Approximation (SAA) [6], we can approximate the solution of (2) by

$$A := \min_{\boldsymbol{z}} \quad \mathbb{E}^{\boldsymbol{y} \sim \mathbb{P}(\boldsymbol{z})} \left[h\left(\boldsymbol{y}, \boldsymbol{z} \right) \right]$$

s.t. $\boldsymbol{z} \in \mathcal{Z}_{\tau}$ (5)
 $\operatorname{CVaR}_{1-\beta}^{\boldsymbol{y} \sim \mathbb{P}(\boldsymbol{z})} \left(g(\boldsymbol{y}, \boldsymbol{z}) \right) \leq 0$,

with

 \hat{J}^{SA}

$$\hat{\mathbb{P}}(oldsymbol{z}) := rac{1}{N} \sum_{i=1}^N \delta_{oldsymbol{\hat{y}}^i(oldsymbol{z})}$$

the empirical distribution of $\hat{y}^{i}(z)$, $i=1,\ldots,N$, as in (4).

B. The distributionally robust formulation

While simple to implement, (5) can lead to poor out-ofsample performance and cannot provide tight probabilistic guarantees if we only have access to a limited amount of data or when there is a mismatch between the approximate and true dynamics. Thus, we resort to a distributionally robust formulation of (2) to robustify against both the uncertainty related to the approximate dynamics and the disturbance acting on the system by optimizing the worst-case expectation within an ambiguity set defined using the Wasserstein metric. While our work generalizes to the so-called p-Wasserstein metric with an arbitrary norm, we restrict our attention to the 1-Wasserstein metric with the Euclidean norm.

Definition 3 (Wasserstein metric [19], [20]). Consider distributions $\mathbb{Q}_1, \mathbb{Q}_2 \in \mathcal{M}(\mathcal{Y})$, where $\mathcal{M}(\mathcal{Y})$ is the set of all probability distributions \mathbb{Q} supported on $\mathcal{Y} \subseteq \mathbb{R}^{nT}$ such that $\mathbb{E}[\|\boldsymbol{y}\|] < \infty$. The Wasserstein metric $d_W : \mathcal{M}(\mathcal{Y}) \times \mathcal{M}(\mathcal{Y}) \to \mathbb{R}_{\geq 0}$ between the distributions \mathbb{Q}_1 and \mathbb{Q}_2 is defined as

$$d_W\left(\mathbb{Q}_1,\mathbb{Q}_2\right) := \inf\left\{\int_{\mathcal{Y}^2} \|\boldsymbol{y}_1 - \boldsymbol{y}_2\| \Pi\left(\mathrm{d}\boldsymbol{y}_1,\mathrm{d}\boldsymbol{y}_2\right)\right\},\qquad(6)$$

where Π is the joint distribution of y_1 and y_2 with marginals \mathbb{Q}_1 and \mathbb{Q}_2 .

The distributionally robust (DR) FHOC problem can now be written as

$$\hat{I}^{DR} := \min_{\boldsymbol{z}} \sup_{\mathbb{Q} \in \mathcal{B}^{\varepsilon}(\hat{\mathbb{P}}(\boldsymbol{z}))} \mathbb{E}^{\boldsymbol{y} \sim \mathbb{Q}} \left[h\left(\boldsymbol{y}, \boldsymbol{z}\right) \right] \\
\text{s.t.} \quad \boldsymbol{z} \in \mathcal{Z}_{\tau} \qquad (7) \\
\sup_{\mathbb{Q} \in \mathcal{B}^{\varepsilon}(\hat{\mathbb{P}}(\boldsymbol{z}))} \operatorname{CVaR}_{1-\beta}^{\boldsymbol{y} \sim \mathbb{Q}} \left(g(\boldsymbol{y}, \boldsymbol{z}) \right) \leq 0 ,$$

where the ambiguity set

$$\mathcal{B}^{arepsilon}\left(\hat{\mathbb{P}}(oldsymbol{z})
ight):=\left\{\mathbb{Q}\in\mathcal{M}(\mathcal{Y})\Big|d_{W}\left(\hat{\mathbb{P}}(oldsymbol{z}),\mathbb{Q}
ight)\leqarepsilon
ight\}$$

is defined as a ball of radius ε with respect to the Wasserstein metric (6), centered on the empirical distribution $\hat{\mathbb{P}}(z)$. Differently from other approaches, e.g. [13], [14], we define the center of the ambiguity set as the empirical distribution of the *T*-step predictions $\hat{y}^i(z)$, i = 1, ..., N. As a consequence, the ambiguity set not only depends on the offline collected trajectories and on the approximate predictor, but also on the optimization variable z, i.e., on the initial condition x_0 and the control inputs sequence u. Notice that separately formulating the worst-case in the cost and in the constraints might introduce some conservatism as the solution robustifies against two, possibly different, worst-case distributions.

III. MAIN RESULTS/FINITE SAMPLE GUARANTEES

We now want to derive a lower bound on the radius ε that guarantees, with high probability, that the true distribution $\mathbb{P}_{y|z}$ is contained in the ambiguity set $\mathcal{B}^{\varepsilon}(\hat{\mathbb{P}}(z))$. This will allow us to show that if the DR Problem (7) is feasible and the minimizer z^* attains an optimal cost $\hat{J}^{DR}(z^*)$, then, with high confidence, z^* is a feasible solution to the original Problem (2) and the resulting cost $\bar{J}(z^*)$ is upper bounded by $\hat{J}^{DR}(z^*)$.

We require a technical assumption on the distribution of the multi-step noise $\boldsymbol{\xi} = \bar{H} \boldsymbol{w}$.

Assumption 4 (Light-tail assumption). For some constants a > 1 and $\gamma > 0$

$$\mathcal{E}_{a,b} = \mathbb{E}\left[e^{b\|\boldsymbol{\xi}\|^a}\right] < +\infty$$
.

This assumption is a condition on the decay rate of the tail of the probability distribution \mathbb{P}_w and it is trivially satisfied when $w \sim \mathbb{P}_w$ is sub-Gaussian or when \mathcal{W} is compact.

To help us in the derivation of the upcoming proofs, we define the distribution

$$ar{\mathbb{P}}(oldsymbol{z}) := rac{1}{N} \sum_{i=1}^N \delta_{oldsymbol{ar{y}}^i(oldsymbol{z})} \;,$$

where

$$\begin{split} \bar{\boldsymbol{y}}^{i}(\boldsymbol{z}) &:= \bar{L}\boldsymbol{z} + \bar{\boldsymbol{\xi}}^{i}, \quad \text{for } i = 1, \dots, N \quad, \\ \bar{\boldsymbol{\xi}}^{i} &:= \boldsymbol{y}^{i} - \bar{L}\boldsymbol{z}^{i} \quad, \quad \text{for } i = 1, \dots, N \;. \end{split}$$

For given initial condition and control input sequence z, the future state prediction $\bar{y}^i(z)$, i = 1, ..., N, is computed using the true (unknown) multi-step dynamics \bar{L} and the true multi-step residuals $\bar{\xi}^i$, i = 1, ..., N. Since we employed here the true dynamics to derive the residuals from the available data and to compute the prediction, $\bar{\mathbb{P}}(z)$ is an empirical sampled version of the true distribution $\mathbb{P}_{y|z}$. Note that we introduced these objects for analysis purposes only and, since \bar{L} is unknown, we cannot compute them.

To obtain the ambiguity set radius ε , we start by upper bounding the Wasserstein distance $d_W(\hat{\mathbb{P}}(z), \mathbb{P}_{y|z})$ between the empirical distribution obtained with the approximate dynamics $\hat{\mathbb{P}}(z)$ and the true unknown distribution $\mathbb{P}_{y|z}$, inspired by [21]. We apply the triangle inequality to obtain

$$d_W\left(\hat{\mathbb{P}}(\boldsymbol{z}), \mathbb{P}_{\boldsymbol{y}|\boldsymbol{z}}\right) \leq d_W\left(\hat{\mathbb{P}}(\boldsymbol{z}), \bar{\mathbb{P}}(\boldsymbol{z})\right) + d_W\left(\bar{\mathbb{P}}(\boldsymbol{z}), \mathbb{P}_{\boldsymbol{y}|\boldsymbol{z}}\right).$$

With the following lemmas we upper bound, uniformly in z, the terms appearing on the right-hand side.

Lemma 5. Under Assumption 2, let α specify a risk level with $\alpha \in (0, 1)$. Then, for each $z \in \mathbb{Z}_{\tau}$

$$\mathbb{P}^{\mathcal{ID}}\left\{d_W\left(\hat{\mathbb{P}}(\boldsymbol{z}), \bar{\mathbb{P}}(\boldsymbol{z})\right) \geq \gamma\left(\frac{\alpha}{2}\right) \frac{1}{N} \sum_{i=1}^N \left\|\boldsymbol{z} - \boldsymbol{z}^i\right\|\right\} \leq \frac{\alpha}{2}$$

Proof.

$$egin{aligned} &d_W\left(\hat{\mathbb{P}}(oldsymbol{z}), ar{\mathbb{P}}(oldsymbol{z})
ight) &\leq rac{1}{N}\sum_{i=1}^N \left\| oldsymbol{\hat{y}}^i(oldsymbol{z}) - oldsymbol{ar{y}}^i(oldsymbol{z})
ight\| \ &= rac{1}{N}\sum_{i=1}^N \left\| ig(\hat{L}oldsymbol{z} + oldsymbol{\hat{\xi}}^iig) - ig(oldsymbol{L}oldsymbol{z} + oldsymbol{ar{\xi}}^iig)
ight\| \ &= rac{1}{N}\sum_{i=1}^N \left\| ig(\hat{L}oldsymbol{z} - oldsymbol{ar{L}}oldsymbol{z} + oldsymbol{ar{z}}(oldsymbol{z}) + ig(oldsymbol{ar{z}} - oldsymbol{ar{z}}oldsymbol{z}^iig)
ight\| \ &\leq rac{1}{N}\sum_{i=1}^N \left(\left\| \hat{L} - oldsymbol{ar{L}}
ight\| \|oldsymbol{z} - oldsymbol{z}^iig\|
ight) \ , \end{aligned}$$

where the first step follows from the definition of Wasserstein distance and the last step from the fact that the matrix norm is consistent with the vector norm that induced it. Therefore,

$$\mathbb{P}^{\mathcal{ID}}\left\{d_{W}\left(\hat{\mathbb{P}}(\boldsymbol{z}), \bar{\mathbb{P}}(\boldsymbol{z})\right) \geq \gamma\left(\frac{\alpha}{2}\right)\left(\frac{1}{N}\sum_{i=1}^{N}\left\|\boldsymbol{z}-\boldsymbol{z}^{i}\right\|\right)\right\}$$
$$\leq \mathbb{P}^{\mathcal{ID}}\left\{\left\|\hat{L}-\bar{L}\right\|\frac{1}{N}\sum_{i=1}^{N}\left\|\boldsymbol{z}-\boldsymbol{z}^{i}\right\| \geq \gamma\left(\frac{\alpha}{2}\right)\frac{1}{N}\sum_{i=1}^{N}\left\|\boldsymbol{z}-\boldsymbol{z}^{i}\right\|\right\} \leq \frac{\alpha}{2}.$$

The derived bound on the distance $d_W(\hat{\mathbb{P}}(z), \bar{\mathbb{P}}(z))$ depends on the model mismatch $\|\hat{L} - \bar{L}\|$ and on the average distance between the z and the collected z^i , i = 1, ..., N.

Lemma 6. ([22, Theorem 2]): Under Assumption 4, for nT > 2, for all $\kappa > 0$, $N \in \mathbb{N}$, and $z \in \mathbb{Z}_{\tau}$

$$\mathbb{P}^{N}\left\{d_{W}\left(\bar{\mathbb{P}}(\boldsymbol{z}), \mathbb{P}_{\boldsymbol{y}|\boldsymbol{z}}\right) \geq \kappa\right\} \leq \begin{cases} c_{1} \exp\left(-c_{2}N\kappa^{nT}\right) & \text{if } \kappa \leq 1\\ c_{1} \exp\left(-c_{2}N\kappa^{a}\right) & \text{if } \kappa > 1 \end{cases},$$

where \mathbb{P}^N is the N-fold distribution of the data generating process (1). The positive constants c_1 and c_2 depend on the dimensions of $\boldsymbol{\xi}$ and on the constants a, b and $\mathcal{E}_{a,b}$.

We are now in the position to prove a finite sample guarantee result.

Theorem 7. Under Assumptions 2 and 4, and for a risk level $\alpha \in (0, 1)$, define the radius of the ambiguity set

$$\varepsilon(\alpha, \boldsymbol{z}) := \varepsilon_1(\alpha) \left(\frac{1}{N} \sum_{i=1}^N \left\| \boldsymbol{z} - \boldsymbol{z}^i \right\| \right) + \varepsilon_2(\alpha) , \qquad (8)$$

with

$$\varepsilon_{1}(\alpha) := \gamma\left(\frac{\alpha}{2}\right), \quad \varepsilon_{2}(\alpha) := \begin{cases} \left(\frac{\log\left(c_{1}\frac{2}{\alpha}\right)}{c_{2}N}\right)^{\frac{1}{nT}} \text{ if } N \ge \frac{\log\left(c_{1}\frac{2}{\alpha}\right)}{c_{2}N} \\ \left(\frac{\log\left(c_{1}\frac{2}{\alpha}\right)}{c_{2}N}\right)^{\frac{1}{a}} \text{ if } N < \frac{\log\left(c_{1}\frac{2}{\alpha}\right)}{c_{2}N} \end{cases}.$$

Let $\hat{J}^{DR}(\boldsymbol{z}^*)$ and \boldsymbol{z}^* be the optimal value and a feasible optimizer of the distributionally robust FHOC Problem (7) with decision-dependent ambiguity set $\mathcal{B}^{\varepsilon}(\hat{\mathbb{P}}(\boldsymbol{z}))$ of radius $\varepsilon = \varepsilon(\alpha, \boldsymbol{z})$ as defined in (8). Then,

$$\begin{cases} \mathbb{P}\left\{\bar{J}(\boldsymbol{z}^{*}) \leq \hat{J}^{DR}\left(\boldsymbol{z}^{*}\right)\right\} \geq 1 - \alpha \\ \mathbb{P}\left\{C VaR_{1-\beta}^{\boldsymbol{y} \sim \mathbb{P}_{\boldsymbol{y} \mid \boldsymbol{z}^{*}}}\left(g(\boldsymbol{y}, \boldsymbol{z}^{*})\right) \leq 0\right\} \geq 1 - \alpha \\ where \ \mathbb{P} = \mathbb{P}^{\mathcal{ID}} \times \mathbb{P}^{N}. \end{cases}$$

Proof. The claim is a direct consequence of Lemma 6 and Lemma 5. Choosing $\varepsilon = \varepsilon(\alpha, z)$ as in (8) ensures that

$$\mathbb{P}\left\{d_{W}\left(\hat{\mathbb{P}}(\boldsymbol{z}), \mathbb{P}_{\boldsymbol{y}|\boldsymbol{z}}\right) \geq \varepsilon(\alpha, \boldsymbol{z})\right\} \leq \\ \leq \mathbb{P}^{\mathcal{ID}}\left\{d_{W}\left(\hat{\mathbb{P}}(\boldsymbol{z}), \bar{\mathbb{P}}(\boldsymbol{z})\right) \geq \varepsilon_{1}(\alpha)\frac{1}{N}\sum_{i=1}^{N}\left\|\boldsymbol{z} - \boldsymbol{z}^{i}\right\|\right\} + \\ + \mathbb{P}^{N}\left\{d_{W}\left(\bar{\mathbb{P}}(\boldsymbol{z}), \mathbb{P}_{\boldsymbol{y}|\boldsymbol{z}}\right) \geq \varepsilon_{2}(\alpha)\right\} \\ \leq \frac{\alpha}{2} + \frac{\alpha}{2} = \alpha \ .$$

The claim follows from the definition of the ambiguity set $\mathcal{B}^{\varepsilon}(\hat{\mathbb{P}}(z))$ and of the DR Problem (7).

The radius in (8) is the sum of two components, one related to limited number of available samples through Lemma 6 and the other to the error in the system dynamics through Lemma 5. While the derivation of $\varepsilon(\alpha, z)$ helps us derive theoretical guarantees, it typically leads to conservative bounds in practice. We investigate a practical data-driven approach for choosing the radius $\varepsilon(\alpha, z)$ in Section IV.

The supremums over probability distributions appearing in the DR Problem (7) make it an infinite-dimensional optimization program. In the following theorem, leveraging recent results in the DR optimization literature, we show that, for the class of problems with convex piecewise affine cost and constraint functions, Problem (7) can be reformulated as a tractable, finite-dimensional convex program. **Theorem 8.** Let Assumption 2 and Assumption 4 hold. Let $\mathbf{y} \in \mathbb{R}^{nT}$, consider a cost function $h(\mathbf{y}, \mathbf{z}) = \max_{j \leq N_j} h_j(\mathbf{y}, \mathbf{z})$, with $h_j(\mathbf{y}, \mathbf{z}) = a_j \mathbf{y} + b_j \mathbf{z} + c_j$ and a constraint function $g(\mathbf{y}, \mathbf{z}) = \max_{k \leq N_k} g_k(\mathbf{y}, \mathbf{z})$, with $g_k(\mathbf{y}, \mathbf{z}) = d_k \mathbf{y} + e_k \mathbf{z} + f_k$. Then, the DR FHOC Problem (7) can be formulated as

$$\hat{J}^{DR} = \inf_{\boldsymbol{z}, s_i} \underline{\lambda} \left(\varepsilon_1(\alpha) \frac{1}{N} \sum_{i=1}^N \|\boldsymbol{z} - \boldsymbol{z}^i\| + \varepsilon_2(\alpha) \right) + \frac{1}{N} \sum_{i=1}^N s_i$$

s.t. $\boldsymbol{z} \in \mathcal{Z}_{CVaR}$
 $a_j (\hat{L} \boldsymbol{z} + \hat{\boldsymbol{\xi}}^i) + b_j \boldsymbol{z} + c_j \leq s_i$
 $\forall i = 1, \dots, N, \ j = 1, \dots, N_j$,

with $Z_{CVaR} =$

$$\begin{cases} \mathbf{z} \in \mathcal{Z}_{\tau} \begin{vmatrix} \exists \ t, q_i \ such \ that \\ \underline{\theta} \left(\varepsilon_1(\alpha) \frac{1}{N} \sum_{i=1}^N \left\| \mathbf{z} - \mathbf{z}^i \right\| + \varepsilon_2(\alpha) \right) + \frac{1}{N} \sum_{i=1}^N q_i \le t\beta \\ \left(d_k \left(\hat{L} \mathbf{z} + \hat{\mathbf{\xi}}^i \right) + e_k \mathbf{z} + f_k + t \right)_+ \le q_i \\ \forall \ i = 1, \dots, N, \ k = 1, \dots, N_k \end{vmatrix} \right\} , \\ \underline{\lambda} = \max_{j \le N_j} \|a_j\|, \quad \underline{\theta} = \max_{k \le N_k} \|d_k\|. \end{cases}$$

Proof. For a convex piecewise affine cost function h(y, z), we can apply Corollary 5.1 of [20] to reformulate the supremum appearing in the worst-case expectation of (7) as

$$\sup_{\mathbb{Q}\in\mathcal{B}^{\epsilon}\left(\hat{\mathbb{P}}(\boldsymbol{z})\right)} \mathbb{E}^{\boldsymbol{y}\sim\mathbb{Q}}\left[h\left(\boldsymbol{y},\boldsymbol{z}\right)\right] = \inf_{\lambda,s_{i}} \lambda \varepsilon(\alpha,\boldsymbol{z}) + \frac{1}{N} \sum_{i=1}^{N} s_{i}$$

s.t. $a_{j}\hat{\boldsymbol{y}}^{i}(\boldsymbol{z}) + b_{j}\boldsymbol{z} + c_{j} \leq s_{i}$ (9)
 $\|a_{j}\| \leq \lambda$
 $\forall i=1,\ldots,N, \ j=1,\ldots,N_{i}$.

From the definition of CVaR, the constraints in Problem (7) can be written as the set

$$\begin{split} \mathcal{Z}_{\text{CVaR}} &= \left\{ \left. \boldsymbol{z} \in \mathcal{Z}_{\tau} \right| \sup_{\mathbb{Q} \in \mathcal{B}^{\varepsilon}\left(\hat{\mathbb{P}}(\boldsymbol{z})\right)} \text{CVaR}_{1-\beta}^{\boldsymbol{y} \sim \mathbb{Q}}\left(g(\boldsymbol{y}, \boldsymbol{z})\right) \leq 0 \right\} \\ &= \left\{ \left. \boldsymbol{z} \in \mathcal{Z}_{\tau} \right| \sup_{\mathbb{Q} \in \mathcal{B}^{\varepsilon}\left(\hat{\mathbb{P}}(\boldsymbol{z})\right)} \inf_{t \in \mathbb{R}} \left[\mathbb{E}^{\boldsymbol{y} \sim \mathbb{Q}}\left[(g\left(\boldsymbol{y}, \boldsymbol{z}\right) + t)_{+}\right] - t\beta \right] \leq 0 \right\}. \end{split}$$

For a convex piecewise affine constraint function g(y, z), following Proposition V.1 of [23], we can reformulate the supremum appearing in the constraint as

$$\mathcal{Z}_{\text{CVaR}} = \left\{ \boldsymbol{z} \in \mathcal{Z}_{\tau} \left| \begin{array}{l} \exists t, \theta, q_i \text{ such that} \\ \theta \varepsilon(\alpha, \boldsymbol{z}) + \frac{1}{N} \sum_{i=1}^{N} q_i \leq t\beta \\ \left(d_k \hat{\boldsymbol{y}}^i(\boldsymbol{z}) + e_k \boldsymbol{z} + f_k + t \right)_+ \leq q_i \\ \|d_k\| \leq \theta \\ \forall i = 1, \dots, N, \ k = 1, \dots, K_k \end{array} \right\}$$
(10)

Substituting (4) in (9) and (10), and taking the infimum over z results in the finite-dimensional convex program

$$\begin{split} \hat{J}^{DR} &= \inf_{\boldsymbol{z}, s_i} \underline{\lambda} \varepsilon(\alpha, \boldsymbol{z}) + \frac{1}{N} \sum_{i=1}^{N} s_i \\ \text{s.t. } \boldsymbol{z} \in \mathcal{Z}_{\text{CVaR}} \\ & a_j \left(\hat{L} \boldsymbol{z} + \hat{\boldsymbol{\xi}}^i \right) + b_j \boldsymbol{z} + c_j \leq s_i \\ & \underline{\lambda} = \max_{j \leq N_j} \|a_j\| \quad \forall \ i = 1, \dots, N, \ j = 1, \dots, N_j \end{split}$$

with

$$\mathcal{Z}_{\text{CVaR}} = \left\{ \boldsymbol{z} \in \mathcal{Z}_{\tau} \middle| \begin{array}{l} \exists t, q_i \text{ such that} \\ \underline{\theta} \varepsilon(\alpha, \boldsymbol{z}) + \frac{1}{N} \sum_{i=1}^{N} q_i \leq t\beta \\ \left(d_k \left(\hat{L} \boldsymbol{z} + \hat{\boldsymbol{\xi}}^i \right) + e_k \boldsymbol{z} + f_k + t \right)_+ \leq q_i \\ \underline{\theta} = \max_{k \leq N_k} \|d_k\| \\ \forall i = 1, \dots, N, \ k = 1, \dots, N_k \end{array} \right\}.$$

Since $\boldsymbol{y} \in \mathbb{R}^{nT}$, $\underline{\lambda}$ and $\underline{\theta}$ are not optimization variable, leading to an exact convex reformulation. As the bound on the radius $\varepsilon(\alpha, \boldsymbol{z})$ holds uniformly in \boldsymbol{z} , the claim follows.

IV. NUMERICAL EXAMPLE

We consider the system

$$oldsymbol{x}_{k+1} = egin{bmatrix} 0.9 & 0.1 \ 0.05 & 0.9 \end{bmatrix} oldsymbol{x}_k + egin{bmatrix} 0 \ 1 \end{bmatrix} oldsymbol{u}_k + oldsymbol{w}_k$$

with additive disturbance $\boldsymbol{w}_k \sim \mathcal{N}(0, 0.03^2)$. We assume that we have access to a dataset $\mathcal{D}^{N,T}$ comprising N trajectories of length T = 5, collected by applying a random input sequence $\boldsymbol{u}_k^i \sim \mathcal{N}(0, 0.5^2)$, $k = 1, \ldots, T$, $i = 1, \ldots, N$, starting from random initial conditions $\boldsymbol{x}_0^i \sim \mathcal{N}(0, 0.5^2)$, $i = 1, \ldots, N$. We consider cost and constraint functions

$$h(\mathbf{y}, \mathbf{z}) = \|\mathbf{y}_{(1)} - \mathbb{1}\|_1,$$
 (11)

$$g(\boldsymbol{y}, \boldsymbol{z}) = \max\{[(\boldsymbol{y}_{(1)} - \mathbb{1})^{\top}, \boldsymbol{y}_{(2)}^{\top}]^{\top}\}, \qquad (12)$$

where $y_{(1)}$ and $y_{(2)}$ represent the first and the second entry of the predicted future states, $\|\cdot\|_1$ is the 1-norm, and $\mathbb{1}$ is a vector of ones of appropriate dimensions. The cost function (11) requires the first entry of the state to stay as close as possible to a constant reference of 1, while the constraint (12) enforces $y_{(1)}$ to be smaller than 1 and $y_{(2)}$ to be positive through a CVaR constraint of level $\beta = 0.2$. To guarantee the existence of a feasible solution at each MPC step, we consider a slack formulation for the CVaR constraint, with the slack variable linearly weighted in the cost function with a weight of 10^6 . For all the simulations, the initial condition is set to $x_0 = [0.9, 0.9]^{\top}$ and the MPC horizon is set to 30 steps.

To study the effect of different ambiguity set radius on the closed-loop cost and violations, we first fix a dataset of size of N = 10 and a closed-loop noise realization. Instead of assuming a given \hat{L} , we use the data in $\mathcal{D}^{N,T}$ to compute a T-step linear predictor by solving the following least-squares minimization problem

$$\hat{L} := \arg\min_{L} \quad \sum_{i=1}^{N} \left\| L \boldsymbol{z}^{i} - \boldsymbol{y}^{i} \right\|^{2} , \qquad (13)$$

under the assumption that the collected data is sufficiently informative. To improve the quality of the predictor, we reduce the number of unknowns by enforcing on \hat{L} the block triangular causal structure of \bar{L} . With the identified dynamics \hat{L} we compute the N T-step residuals as in (3). We intentionally choose a low signal-to-noise ratio and a small number of trajectories in the offline-collected dataset to highlight how limited data availability and model mismatch (in this case due to the identification process) can hamper MPC performance. We then proceed to solve



Fig. 1. Closed-loop cost (left) and violations (right) for different values of ε_1 and ε_2 .

the DRMPC for all the possible combinations of $\varepsilon_1, \varepsilon_2 \in \{10^{-7}, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 10^{0}\}$. The resulting closed-loop cost and violations for different values of the ambiguity set radius are shown in Figure 1. For small values of the radius, the controller experiences a large number of constraint violations, as the radius of the ambiguity set is increased, the number of violations is reduced at the expense of higher closed-loop costs.

As mentioned in [16], the theoretical bound on $\gamma(\alpha)$ in Assumption 2 can be quite conservative in practice. Moreover, the use of the triangle inequality in Theorem 7 can introduce some extra conservatism, leading to a loose bound on the radius $\varepsilon(\alpha, z)$. Informed by the expression in (8), we aim to obtain a less conservative ambiguity set radius by directly exploiting data to approximate the Wasserstein distance $d_W(\hat{\mathbb{P}}(z), \mathbb{P}_{y|z})$.

Algorithm 1: Data-driven empirical radius estimate
Input: $\mathcal{D}^{N,T}$.
Output: \hat{L} , ε_1 , ε_2 .
1: for $\ell = 1$ to N do
2: define $\mathcal{D}_{I_{\ell}}^{N-1,T}$ the leave-one-out dataset with indices
$I_{\ell} = \{1, \dots, N\} \setminus \ell;$
3: compute \tilde{L}^{ℓ} from $\mathcal{D}_{I_{\ell}}^{N-1,T}$ by least-squares regression as in (13)
4: compute residuals $\tilde{\boldsymbol{\xi}}_{\ell}^{i} := \boldsymbol{y}^{i} - \tilde{L}_{\ell} \boldsymbol{z}^{i}, i = 1, \dots, N;$
5: compute $V_{\ell} = \frac{1}{N-1} \sum_{i \in I_{\ell}} \ \boldsymbol{z}^{\ell} - \boldsymbol{z}^{i} \ $;
6: compute $E_{\ell} = \frac{1}{N^2 - N} \sum_{i \in I_{\ell}} \left\ \boldsymbol{y}^{\ell} - \tilde{\boldsymbol{y}}^{i}(\boldsymbol{z}^{\ell}) \right\ $ with
$ ilde{oldsymbol{y}}^{i}(oldsymbol{z}^{\ell}) = ilde{L}_{\ell}oldsymbol{z}^{\ell} + ilde{oldsymbol{\xi}}^{i}_{\ell};$
7: end for
8: return $\hat{L} = \frac{1}{N} \sum_{\ell=1}^{N} \tilde{L}_{\ell}$ and $\{\varepsilon_1 \ge 0, \varepsilon_2 \ge 0\}$ that minimize
$\sum_{\ell=1}^{N} \ \varepsilon_1 V_{\ell} + \varepsilon_2 - E_{\ell}\ .$

In Algorithm 1 we use a leave-one-out procedure to obtain an estimate of the dynamics and of the ambiguity set parameters ε_1 and ε_2 . For each leave-one-out dataset, E_ℓ is the Wasserstein distance $d_W(\tilde{\mathbb{P}}_{i\in I_\ell}(\boldsymbol{z}^\ell), \tilde{\mathbb{P}}(\boldsymbol{z}^\ell))$ between the leave-one-out empirical distribution $\tilde{\mathbb{P}}_{i\in I_\ell}(\boldsymbol{z}^\ell) := \frac{1}{N-1} \sum_{i\in I_\ell} \delta_{\tilde{\boldsymbol{y}}^i}(\boldsymbol{z}^\ell)$ and the one that uses all the samples $\tilde{\mathbb{P}}(\boldsymbol{z}) := \frac{1}{N} \sum_{i=1}^N \delta_{\tilde{\boldsymbol{y}}^i}(\boldsymbol{z}^\ell)$ in $\boldsymbol{z} = \boldsymbol{z}^j$. The only difference between the two distributions is that $\tilde{\mathbb{P}}(\boldsymbol{z}^\ell)$ contains one extra impulse in $\tilde{\boldsymbol{y}}^\ell(\boldsymbol{z}^\ell)$. As the predictions are computed for $\boldsymbol{z} = \boldsymbol{z}^\ell$, by construction we have $\tilde{\boldsymbol{y}}^\ell(\boldsymbol{z}^\ell) = \boldsymbol{y}^\ell$ independently of the dynamics, making E_ℓ the distributional

mismatch that we can expect from a new independent observation $(\boldsymbol{z}^{\ell}, \boldsymbol{y}^{\ell}) \sim \mathbb{P}_{\boldsymbol{y}|\boldsymbol{z}=\boldsymbol{z}^{\ell}}$.

In Figure 2 we compare the closed-loop cost and Closed-loop violations for the proposed DR approach with datadriven radius obtained as in Algorithm 1 (DRMPC) and for SAA approach (5) (SAAMPC) that is obtained by setting $\varepsilon(\alpha, z) = 0$. We perform simulations for $N = \{10, 20, 30, 40\}$, each repeated for 50 realizations of the identification dataset and closed-loop noise trajectory. We can observe that the DRMPC approach can reduce the chance of closed-loop constraint violations, with only a limited increase in closed-loop cost. As N increases, the uncertainty decreases, and we observe a reduction in the performance difference between DRMPC and SAAMPC for both closed-loop cost and constraint violations.

V. CONCLUSIONS

We presented a novel data-driven DRMPC formulation for unknown discrete-time linear time-invariant systems affected by additive uncertainties. We obtained finite sample probabilistic guarantees on the worst-case expectation and CVaR constraint in the presence of uncertainty on both the model of the dynamics and on the disturbance distribution. We then derived a finite-dimensional tractable reformulation of the DR problem for convex piecewise affine cost and constraint functions. Finally, we described a simple datadriven algorithm to obtain an empirical ambiguity set radius estimate and tested the proposed DRMPC against the SAA formulation. The numerical simulations demonstrated the effectiveness of the proposed DRMPC, that achieved a strong reduction in the number of closed-loop constraint violations without substantial increase in the attained closed-loop cost, even when very limited information regarding the dynamics and the disturbance is available. Future work focuses on the analysis of the DRMPC closed-loop properties such as stability and recursive feasibility.



Fig. 2. Comparison of Closed-loop cost (left) and Closed-loop violations (right) for dataset sizes $N = \{10, 20, 30, 40\}$ between SAAMPC and DRMPC with radius tuned as in Algorithm 1.

REFERENCES

- D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.
- [2] W. Langson, I. Chryssochoos, S. Raković, and D. Q. Mayne, "Robust model predictive control using tubes," *Automatica*, vol. 40, no. 1, pp. 125–133, 2004.
- [3] D. Munoz-Carpintero, M. Cannon, and B. Kouvaritakis, "Recursively feasible robust mpc for linear systems with additive and multiplicative uncertainty using optimized polytopic dynamics," in 52nd IEEE Conference on Decision and Control. IEEE, 2013, pp. 1101–1106.
- [4] M. Cannon, B. Kouvaritakis, and X. Wu, "Probabilistic constrained mpc for multiplicative and additive stochastic uncertainty," *IEEE Transactions on Automatic Control*, vol. 54, no. 7, pp. 1626–1632, 2009.
- [5] M. Cannon, B. Kouvaritakis, S. V. Raković, and Q. Cheng, "Stochastic tubes in model predictive control with probabilistic constraints," *IEEE Transactions on Automatic Control*, vol. 56, no. 1, pp. 194–200, 2010.
- [6] A. J. Kleywegt, A. Shapiro, and T. Homem-de Mello, "The sample average approximation method for stochastic discrete optimization," *SIAM Journal on Optimization*, vol. 12, no. 2, pp. 479–502, 2002.
- [7] G. C. Calafiore and M. C. Campi, "The scenario approach to robust control design," *IEEE Transactions on automatic control*, vol. 51, no. 5, pp. 742–753, 2006.
- [8] F. Micheli and J. Lygeros, "Scenario-based stochastic mpc for systems with uncertain dynamics," in 2022 European Control Conference (ECC). IEEE, 2022, pp. 833–838.
- [9] G. C. Calafiore and L. Fagiano, "Robust model predictive control via scenario optimization," *IEEE Transactions on Automatic Control*, vol. 58, no. 1, pp. 219–224, 2012.
- [10] B. P. Van Parys, D. Kuhn, P. J. Goulart, and M. Morari, "Distributionally robust control of constrained stochastic systems," *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 430–442, 2015.
- [11] P. Coppens and P. Patrinos, "Data-driven distributionally robust mpc for constrained stochastic systems," *IEEE Control Systems Letters*, vol. 6, pp. 1274–1279, 2021.
- [12] B. Li, Y. Tan, A.-G. Wu, and G.-R. Duan, "A distributionally robust optimization based method for stochastic model predictive control," *IEEE Transactions on Automatic Control*, 2021.
- [13] C. Mark and S. Liu, "Stochastic mpc with distributionally robust chance constraints," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 7136– 7141, 2020.
- [14] Z. Zhong, E. A. del Rio-Chanona, and P. Petsagkourakis, "Datadriven distributionally robust mpc using the wasserstein metric," arXiv preprint arXiv:2105.08414, 2021.
- [15] A. Zolanvari and A. Cherukuri, "Data-driven distributionally robust iterative risk-constrained model predictive control," *arXiv preprint arXiv:2111.12977*, 2021.
- [16] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, "On the sample complexity of the linear quadratic regulator," *Foundations of Computational Mathematics*, vol. 20, no. 4, pp. 633–679, 2020.
- [17] T. Sarkar and A. Rakhlin, "Near optimal finite time identification of arbitrary linear dynamical systems," in *International Conference on Machine Learning*. PMLR, 2019, pp. 5610–5618.
- [18] M. Simchowitz, H. Mania, S. Tu, M. I. Jordan, and B. Recht, "Learning without mixing: Towards a sharp analysis of linear system identification," in *Conference On Learning Theory*. PMLR, 2018, pp. 439–473.
- [19] L. V. Kantorovich and S. Rubinshtein, "On a space of totally additive functions," *Vestnik of the St. Petersburg University: Mathematics*, vol. 13, no. 7, pp. 52–59, 1958.
- [20] P. Mohajerin Esfahani and D. 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.
- [21] R. Kannan, G. Bayraksan, and J. R. Luedtke, "Residuals-based distributionally robust optimization with covariate information," arXiv preprint arXiv:2012.01088, 2020.
- [22] N. Fournier and A. Guillin, "On the rate of convergence in wasserstein distance of the empirical measure," *Probability Theory and Related Fields*, vol. 162, no. 3, pp. 707–738, 2015.
- [23] A. R. Hota, A. Cherukuri, and J. Lygeros, "Data-driven chance constrained optimization under wasserstein ambiguity sets," in 2019 American Control Conference (ACC). IEEE, 2019, pp. 1501–1506.