

The optimal transport paradigm enables data compression in data-driven robust control

Questa è la versione preprint della seguente opera:

Original

The optimal transport paradigm enables data compression in data-driven robust control / Fabiani, Filippo; Goulart, Paul J.. - (2021), pp. 2412-2417. (2021 American Control Conference New Orleans, LA, USA 25-28 May 2021) [10.23919/ACC50511.2021.9483139].

Availability:

This version is available at: 20.500.11771/25777

Publisher:

IEEE

Published

DOI:10.23919/ACC50511.2021.9483139

Terms of use:

This publication is made accessible in accordance with the terms for deposit in the institutional repository, as defined by the IMT School for Advanced Studies Lucca's Open Access Policy. (https://library.imtlucca.it/sites/default/files/regolamento-policy-open-access-imtlib_0.pdf).

Si prega di consultare le pagine informative dell'editore relative alle politiche di autoarchiviazione.

(Article begins on next page)

The optimal transport paradigm enables data compression in data-driven robust control

Filippo Fabiani and Paul J. Goulart

Abstract

A new data-enabled control technique for uncertain linear time-invariant systems, recently conceived by Coulson et al., builds upon the direct optimization of controllers over input/output pairs drawn from a large dataset. We adopt an optimal transport-based method for compressing such large dataset to a smaller synthetic dataset of representative behaviours, aiming to alleviate the computational burden of controllers to be implemented online. Specifically, the synthetic data are determined by minimizing the Wasserstein distance between atomic distributions supported on both the original dataset and the compressed one. We show that a distributionally robust control law computed using the compressed data enjoys the same type of performance guarantees as the original dataset, at the price of enlarging the ambiguity set by an easily computable and well-behaved quantity. Numerical simulations confirm that the control performance with the synthetic data is comparable to the one obtained with the original data, but with significantly less computation required.

I. INTRODUCTION

In recent years, traditional model-based controller design has been giving way to data-driven approaches. The growing complexity of modern control problems, which often precludes the use of classical system identification procedures, along with the increasing data storage capacity, make *learning from data* a timely and attractive new direction for the system-and-control community.

Two recent trends in control originated from different, but nearly complementary, views of a well-known result in subspace identification, i.e., the so called *Willems' lemma* [1]. Specifically, this technical result establishes that a nonparametric linear time-invariant (LTI) realization of an unknown system can be recovered from a data matrix of noise-free input/output measurements. Thus, an early research direction adopted the Willems' lemma to perform explicit data-dependent parametrization analysis and control of systems [2]–[4]. The second more recent direction, which is central in this paper, designed data-enabled (robust) optimal controllers without this intermediate modelling step [5]–[7].

To ensure that the column space of the data matrix spans all possible trajectories of a corresponding LTI system, however, the Willems' lemma relies on the *persistent excitation* of the system input. This directly translates to a requirement on the input/output observations length, and therefore one might be induced to collect extremely large datasets. As an undesired consequence, especially for the second line of research introduced above, this may pose several computational challenges to the real-time implementation of optimization-based control laws, thus limiting the scope of possible applications. Systems that require fast sampling rates, for instance, may not be amenable to this new approach to control design. Moreover, practitioners may choose short control horizons to alleviate the computational burden, thereby resulting in poor control performance or even system instability. Designing suitable procedures to select the most representative data within large (possibly noisy) datasets is essential to overcome these problems.

Data compression and dimensionality reduction have been widely adopted in the systems-and-control community to extrapolate synthetic sets of samples that “best” capture the information content of noise-corrupted datasets. Statistical and manifold learning [8], [9], principal component analysis [10]–[12], or subspace identification approaches [13]–[15], represent only a few of the most popular techniques. Conceptually, these methods tacitly neglect the stochastic nature of the noise itself, eliminating the less

The authors are with the Department of Engineering Science, University of Oxford, OX1 3PJ, United Kingdom ({filippo.fabiani, paul.goulart}@eng.ox.ac.uk). This work was partially supported through the Governments modern industrial strategy by Innovate UK, part of UK Research and Innovation, under Project LEO (Ref. 104781).

significant components (according to some criterion), or those potentially associated with noise, at a later stage only.

In contrast to the aforementioned literature, we propose to design a synthetic set of samples via an offline procedure based on optimal transport [16], [17] techniques. Specifically, the *Wasserstein distance* represents the key tool to define an optimization problem that minimizes the distance between the discrete measure associated with our synthetic dataset and the empirical distribution of the original noisy data. The resulting linear program (LP) turns out to be a variational Wasserstein problem [17, §9], wherein each atom of the synthetic dataset can be identified with a specific barycentre for a cluster of the original samples, thereby serving as a proxy a whole set of observed system behaviours. In the context of the data-enabled robust control technique of [6], we show that a control law computed via the repeated solution of a distributionally robust optimization problem, built upon a Wasserstein ball as an ambiguity set, enjoys exactly the same type of performance guarantees as a controller designed using the full original dataset, at the price of enlarging the ambiguity radius by a quantity that depends on the number of synthetic atoms adopted. We show that this additional ambiguity radius vanishes as the number of atoms in our synthetic dataset tends to the cardinality of the original dataset.

The paper is organized as follows: we recall some fundamentals of optimal transport in §II, and introduce the robust control problem addressed in §III. We formalize and discuss the variational Wasserstein problem associated with the design of synthetic datasets in §IV. Finally, we compare the control performances obtained over the original and compressed datasets through numerical simulations in §V.

Notation: For vectors $(x_1, \dots, x_N) \in \mathbb{R}^n$ and $\mathcal{I} := \{1, \dots, N\}$, we denote $\text{col}((x_i)_{i \in \mathcal{I}}) := (x_1^\top, \dots, x_N^\top)^\top$. Given a matrix $X \in \mathbb{R}^{n \times m}$, its (i, j) entry is denoted by $[X]_{i,j}$. The symbol $\langle \cdot, \cdot \rangle$ denotes an inner product in the appropriate space, i.e., $\langle x, y \rangle = x^\top y$ for $(x, y) \in \mathbb{R}^n$ and $\langle X, Y \rangle = \text{trace}(X^\top Y)$ for $(X, Y) \in \mathbb{R}^{n \times m}$. The probability simplex $\Sigma_n := \{\sigma \in \mathbb{R}_+^n \mid \mathbf{1}_n^\top \sigma = 1\}$, where $\mathbf{1}_n$ is n -dimensional vector of elements equal to 1. For any point $x \in \mathbb{R}^n$, δ_x is the Dirac unit mass on x . Given a collection of points $\{x_i\}_{i \in \mathcal{I}} \in \mathbb{R}^n$, $\text{conv}(\{x_i\}_{i \in \mathcal{I}})$ represents their convex hull, while $\hat{\mathbb{P}} = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \delta_{x_i}$ is the associated uniform discrete probability distribution. The dual norm of an arbitrary norm $\|\cdot\|$ on \mathbb{R}^n is $\|x\|_* := \sup_{\|y\| \leq 1} \langle x, y \rangle$. The conjugate function of $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined by $f^*(\xi) := \sup_{x \in \mathbb{R}^n} \langle \xi, x \rangle - f(x)$. Quantities denoted with $(\hat{\cdot})$ are either directly measured, or depend on data.

II. FUNDAMENTALS OF OPTIMAL TRANSPORT

We first recall the definition of the Wasserstein distance for continuous measures. Then, by considering discrete distributions, we define the associated optimal transport problem that we will apply to our synthetic and empirical datasets.

A. Wasserstein distance between probability measures

Let Ω be an arbitrary space endowed with a metric d , and $\mathcal{P}(\Omega)$ be the set of Borel probability measures on Ω .

Definition 1: ([16, Ch. 7]) Given any $p \in [1, +\infty)$, the p -Wasserstein distance $\mathcal{W}_p : \mathcal{P}(\Omega) \times \mathcal{P}(\Omega) \rightarrow \mathbb{R}_+$ between two probability measures $(\mathbb{P}, \mathbb{Q}) \in \mathcal{P}(\Omega)$ is defined as

$$\mathcal{W}_p(\mathbb{P}, \mathbb{Q}) := \left(\inf_{\pi \in \Pi(\mathbb{P}, \mathbb{Q})} \int_{\Omega \times \Omega} d^p(x, y) d\pi(x, y) \right)^{1/p}, \quad (1)$$

where $\Pi(\mathbb{P}, \mathbb{Q})$ denotes the set of all probability measures on $\Omega \times \Omega$ that have marginals \mathbb{P} and \mathbb{Q} , respectively. \square

Roughly speaking, the decision variable π of the infinite-dimensional optimization problem in (1) coincides with a transportation plan for moving a mass distribution described by \mathbb{P} to another one described by \mathbb{Q} , while d is the associated transportation cost. Then, given any $\varepsilon > 0$, we define the *Wasserstein ball* of radius ε , centred around the distribution \mathbb{P} , as $\mathcal{B}_\varepsilon(\mathbb{P}) := \{\mathbb{Q} \in \mathcal{P}(\Omega) \mid \mathcal{W}_p(\mathbb{P}, \mathbb{Q}) \leq \varepsilon\}$.

B. Discrete probability distributions

Now, let us consider two families of N and M points in Ω , i.e., $X = \{x_1, \dots, x_N\}$ and $Y = \{y_1, \dots, y_M\}$, respectively. Given weights $\alpha \in \Sigma_N$, $\beta \in \Sigma_M$, we can construct discrete probability distributions $\hat{\mathbb{P}}$ and $\hat{\mathbb{Q}}$ as $\hat{\mathbb{P}} = \sum_{i \in \mathcal{N}} \alpha_i \delta_{x_i}$ and $\hat{\mathbb{Q}} = \sum_{i \in \mathcal{M}} \beta_i \delta_{y_i}$, where $\mathcal{N} := \{1, \dots, N\}$, $\mathcal{M} := \{1, \dots, M\}$. In this special case, the Wasserstein distance happens to correspond to the optimal value of a network problem, as (1) translates to the following LP [17]:

$$\mathcal{W}_p(\hat{\mathbb{P}}, \hat{\mathbb{Q}}) = \min_{T \in \mathcal{T}(\alpha, \beta)} \langle T, D(X, Y) \rangle. \quad (2)$$

Here, $D \in \mathbb{R}^{N \times M}$ is the matrix of pairwise distances between points in X and Y , raised to the power p , defined as $[D]_{i,j} := d^p(x_i, y_j)$, for all $x_i \in X$ and $y_j \in Y$. Moreover, every element of the decision matrix $T \in \mathbb{R}^{N \times M}$ in (2), i.e., $[T]_{i,j} := t_{i,j}$, determines the coupling between pairs $(x_i, y_j) \in X \times Y$, whose value specifies the amount of mass flowing from the point $x_i \in X$ toward the point $y_j \in Y$. For any $\alpha \in \Sigma_N$ and $\beta \in \Sigma_M$, the admissible couplings lie in the feasible set $\mathcal{T}(\alpha, \beta)$, defined as

$$\mathcal{T}(\alpha, \beta) := \{T \in \mathbb{R}_+^{N \times M} \mid T \mathbf{1}_M = \alpha, T^\top \mathbf{1}_N = \beta\}. \quad (3)$$

The set of matrices in (3) is called the *transportation polytope*, since it is convex, bounded and defined by a set of $N + M$ equality constraints. Note that as long as $\mathcal{T}(\alpha, \beta)$ is nonempty, the solution to the LP in (2), attained on the vertices of \mathcal{T} , may not be unique. Finally, since \mathcal{W}_p defines a metric, in the discrete setting (2) we have $\mathcal{W}_p(\hat{\mathbb{P}}, \hat{\mathbb{Q}}) = 0$ if and only if $\alpha = \beta$ [17, Prop. 2.2], and therefore the triangle inequality holds as $\mathcal{W}_p(\hat{\mathbb{P}}, \hat{\mathbb{Q}}) \leq \mathcal{W}_p(\hat{\mathbb{P}}, \hat{\mathbb{G}}) + \mathcal{W}_p(\hat{\mathbb{G}}, \hat{\mathbb{Q}})$, for any discrete measures $\hat{\mathbb{P}}, \hat{\mathbb{Q}}, \hat{\mathbb{G}} \in \mathcal{P}(\Omega)$.

We henceforward focus on the *Kantorovich-Rubinstein distance* obtained by setting $p = 1$, and consequently write $\mathcal{W}(\mathbb{P}, \mathbb{Q})$ without subscript. We will assume that the metric d is induced by an arbitrary norm $\|\cdot\|$ on \mathbb{R}^n .

III. THE DATA-ENABLED CONTROL PARADIGM

We start by formalizing the optimal control problem we wish to consider. After a brief digression on nonparametric models for deterministic LTI systems, we will recall from [6] a distributionally robust reformulation of the control problem.

A. Constrained optimal control of uncertain LTI systems

We consider discrete time, stochastic systems in the form:

$$\begin{cases} x(k+1) = Ax(k) + Bu(k) + E\nu(k), \\ y(k) = Cx(k) + Du(k) + F\nu(k), \end{cases} \quad (4)$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $E \in \mathbb{R}^{n \times q}$, $C \in \mathbb{R}^{\ell \times n}$, $D \in \mathbb{R}^{\ell \times m}$ and $F \in \mathbb{R}^{\ell \times q}$. The state, control input, output and disturbance at time instant $k \in \mathbb{Z}$ are $x(k) \in \mathbb{R}^n$, $u(k) \in \mathbb{R}^m$, $y(k) \in \mathbb{R}^\ell$ and $\nu(k) \in \mathbb{R}^q$, respectively. The uncertainty $\nu(k)$ is drawn from an unknown probability distribution \mathbb{P}_ν , supported on $\Upsilon \subseteq \mathbb{R}^q$. We will assume throughout that the system matrices defining (4) are unknown, and that we have access to input/output measurements only, i.e., $(\hat{u}(k), \hat{y}(k))$, $k \in \mathbb{Z}$. Note that, in view of the dynamics in (4), any output measurement $\hat{y}(k)$ can be affected by the realization of the stochastic disturbance $\nu(k)$, for any $k \in \mathbb{Z}$.

A typical approach to steer the behaviour of (4), particularly in the presence of state or input constraints, is stochastic model predictive control [18]. To this end, we consider a finite horizon control problem over horizon length $K \in \mathbb{N}$, where we aim to design a constrained sequence of control inputs, i.e., $u := \text{col}(u(k), \dots, u(k+K-1)) \in \mathcal{U}$, for some compact, convex set $\mathcal{U} \subseteq \mathbb{R}^{mK}$, while minimizing

a predefined cost function $J : \mathbb{R}^{mK} \times \mathbb{R}^{\ell K} \rightarrow \mathbb{R}$. Specifically, in view of the uncertain nature of (4), the finite horizon control problem translates into the following stochastic program:

$$\inf_{u \in \mathcal{U}} \mathbb{E}_{\mathbb{P}_\nu^K} [J(u, y)], \quad (5)$$

where $\mathbb{P}_\nu^K := \mathbb{P}_\nu \times \dots \times \mathbb{P}_\nu$ is the K -fold product distribution characterizing ν over the whole horizon K . Next, we formulate the same working assumptions as in [6].

Standing Assumption 1: The pair (A, B) is controllable. \square

Standing Assumption 2: For all $(u, y) \in \mathbb{R}^{mK} \times \mathbb{R}^{\ell K}$, $J(u, y)$ is a separable function, namely $J(u, y) := J_1(u) + J_2(y)$, where $J_1 : \mathbb{R}^{mK} \rightarrow \mathbb{R}$, $J_2 : \mathbb{R}^{\ell K} \rightarrow \mathbb{R}$ are convex and continuous. In addition, J_2 is such that $\Xi := \{\xi \in \mathbb{R}^{\ell K} \mid J_2^*(\xi) < \infty\} \subseteq \mathbb{R}^{\ell K}$ is a bounded set. \square

B. Nonparametric models for deterministic LTI systems

Let us first consider a deterministic version of the system in (4), i.e., with $\nu(k) = 0$, for any $k \in \mathbb{Z}$. The Willems' fundamental lemma [1, Th. 1] provides the theoretical means to construct data-consistent, minimal, nonparametric models for unknown, deterministic LTI systems [2]–[4], [6]. Specifically, assume that we have available (noise-free) data from an experiment of length N , $(\{\hat{u}(i)\}_{i=0}^{N-1}, \{\hat{y}(i)\}_{i=0}^{N-1})$, for different time shifts. Without loss of generality, assume also that $k = 0$ corresponds to the instant of initial observation. This data can then be organized within a matrix $\mathcal{H}_K := \text{col}(\hat{\mathcal{U}}_{0,K,N-K+1}, \hat{\mathcal{Y}}_{0,K,N-K+1}) \in \mathbb{R}^{(m+\ell)K \times N-K+1}$, where

$$\hat{\mathcal{U}}_{0,K,N-K+1} := \begin{bmatrix} \hat{u}(0) & \hat{u}(1) & \cdots & \hat{u}(N-K) \\ \hat{u}(1) & \hat{u}(2) & \cdots & \hat{u}(N-K+1) \\ \vdots & \vdots & \ddots & \vdots \\ \hat{u}(K-1) & \hat{u}(K) & \cdots & \hat{u}(N-1) \end{bmatrix} \quad (6)$$

and $\hat{\mathcal{Y}}_{0,K,N-K+1} \in \mathbb{R}^{\ell K \times N-K+1}$ is defined similarly. The first subscript refers to the time index of the top-left entry of a given matrix, the second refers to the number of block-rows, and the third to the number of columns. Note that both $\hat{\mathcal{U}}_{0,K,N-K+1} \in \mathbb{R}^{mK \times N-K+1}$ and $\hat{\mathcal{Y}}_{0,K,N-K+1}$ have constant vector entries along the block anti-diagonals, and therefore \mathcal{H}_K belongs to the class of block-Hankel matrices.

The Willems' fundamental lemma restricts the class of input sequences over the horizon K to the persistently exciting signals, as defined next.

Definition 2: ([1]) A measured signal $\hat{z} \in \mathbb{R}^w$, observed over N samples, is *persistently exciting* of order K if the corresponding Hankel matrix $\hat{\mathcal{Z}}_{0,K,N-K+1}$, defined equivalently to (6), has full rank wK . \square

It follows that for a signal to be persistently exciting of order K , its length N must satisfy $N \geq (w+1)K - 1$. Then, by relying on Definition 2, we restate the Willems' fundamental lemma as follows:

Lemma 1: ([1, Th. 1]) Let $\text{col}(\{\hat{u}(i)\}_{i=0}^{N-1})$ be a persistently exciting control signal of order $n + K$. Then $\text{col}(u, y)$ is a K -long input/output trajectory of the deterministic version of the system in (4) if and only if $\text{col}(u, y) \in \text{Im}(\mathcal{H}_K)$. \square

Lemma 1 establishes that if N is chosen large enough and signals are persistently exciting, then every realisable input/output trajectory of the deterministic system is a linear combination of collected input/output data. In other words, for any K -long input/output trajectory $\text{col}(u, y)$, there will always exist some $g \in \mathbb{R}^{N-K+1}$ such that $\text{col}(u, y) = \mathcal{H}_K g$.

By making use of Lemma 1, our goal is next to restate the deterministic version of the finite horizon control problem in (5) by rearranging a measured N -long, input/output trajectory, $\text{col}(\hat{u}, \hat{y}) := (\{\hat{u}(i)\}_{i=0}^{N-1}, \{\hat{y}(i)\}_{i=0}^{N-1})$. Specifically, as in [6, §III.A], by starting from the current time $k \in \mathbb{N}$, we assume that the control input \hat{u} is persistently exciting of order $K_i + K + n$, for some $K_i \in \mathbb{N}$. We

then split \mathcal{H}_{K_i+K} into block matrices $\hat{\mathcal{U}}_f$, $\hat{\mathcal{Y}}_f$, $\hat{\mathcal{U}}_b$ and $\hat{\mathcal{Y}}_b$. Here, $\hat{\mathcal{U}}_f \in \mathbb{R}^{mK \times N-(K_i+K)+1}$ and $\hat{\mathcal{Y}}_f \in \mathbb{R}^{\ell K \times N-(K_i+K)+1}$ consist of the last K -block rows of \mathcal{H}_{K_i+K} , corresponding to data matrices for the “forward” propagation (i.e., from $k \in \mathbb{N}$ onward) of control sequence and output prediction, respectively. Conversely, $\hat{\mathcal{U}}_b \in \mathbb{R}^{mK_i \times N-(K_i+K)+1}$ and $\hat{\mathcal{Y}}_b \in \mathbb{R}^{\ell K_i \times N-(K_i+K)+1}$, which correspond to the first K_i -block rows of \mathcal{H}_{K_i+K} , define the consistency constraints associated with less recent measurements (i.e., “backward” data), together with $\hat{u}_i := \text{col}(\hat{u}(k-K_i), \dots, \hat{u}(k-1))$ and $\hat{y}_i := \text{col}(\hat{y}(k-K_i), \dots, \hat{y}(k-1))$. Thus, the deterministic version of (5) follows directly from Lemma 1 and reads as:

$$\begin{cases} \min_g & J(\hat{\mathcal{U}}_f g, \hat{\mathcal{Y}}_f g) \\ \text{s.t.} & \begin{bmatrix} \hat{\mathcal{U}}_b \\ \hat{\mathcal{Y}}_b \end{bmatrix} g = \begin{bmatrix} \hat{u}_i \\ \hat{y}_i \end{bmatrix}, \\ & \hat{\mathcal{U}}_f g \in \mathcal{U}, \end{cases} \quad (7)$$

where the decision variable g belongs to $\mathbb{R}^{N-(K_i+K)+1}$.

C. A distributionally robust data-enabled control problem

For uncertain systems, the optimization problem in (7) is complicated by the realization of noise terms $\nu(k)$ drawn from the distribution \mathbb{P}_ν . In particular, system noise complicates satisfaction of the consistency constraint $\hat{\mathcal{Y}}_b g = \hat{y}_i$.

As proposed in [6], a possible approach is to soften this consistency constraint, directly penalizing the term $\hat{\mathcal{Y}}_b g - \hat{y}_i$ in the cost function as follows

$$\min_{g \in \mathcal{G}} J(\hat{\mathcal{U}}_f g, \hat{\mathcal{Y}}_f g) + \rho \|\hat{\mathcal{Y}}_b g - \hat{y}_i\|_1, \quad (8)$$

where $\mathcal{G} := \{g \in \mathbb{R}^{N-(K_i+K)+1} \mid \hat{\mathcal{U}}_f g \in \mathcal{U}, \hat{\mathcal{U}}_b g = \hat{u}_i\}$ depends on input measurements only. The optimization problem in (8) can be manipulated to obtain a distributionally robust, semi-infinite reformulation. Specifically, we note that all random objects can be gathered into a matrix $\begin{bmatrix} \mathcal{Y}_b & y_i \\ \mathcal{U}_f & 0 \end{bmatrix}$, whose j -th row is denoted by κ_j^\top . Any such row corresponds to a random vector distributed according to some probability \mathbb{P}_{κ_j} , and supported on $\Theta_{k_j} \subseteq \mathbb{R}^{N-(K_i+K)+2}$, for all $j \in \{1, \dots, \ell(K_i+K)\}$. Note that every \mathbb{P}_{κ_j} and Θ_{k_j} is determined starting from the unknown distribution \mathbb{P}_ν and support Υ , respectively. Define $\kappa := \text{col}((\kappa_j)_{j=1}^{\ell(K_i+K)})$, a random vector supported on $\Theta := \prod_{j=1}^{\ell(K_i+K)} \Theta_{k_j} \subseteq \mathbb{R}^{\ell(K_i+K)(N-(K_i+K)+2)}$ and distributed according to $\mathbb{P}_\kappa := \prod_{j=1}^{\ell(K_i+K)} \mathbb{P}_{\kappa_j}$, and let $v := \text{col}(g, -1)$ and $\mathcal{V} := \mathcal{G} \times \{-1\} \subseteq \mathbb{R}^{N-(K_i+K)+2}$. With the notation introduced, and taking into account the realization of the random objects in κ , the cost function in (8) turns out to be $J((\hat{\mathcal{U}}_f, 0)v, (\hat{\kappa}_{\ell(K_i+K)+1}v, \dots, \hat{\kappa}_{\ell(K_i+K)}v)) + \rho \|(\hat{\kappa}_1 v, \dots, \hat{\kappa}_{\ell(K_i+K)}v)\|_1 =: f(\hat{\kappa}, v)$. Therefore, based on the measurements $\hat{\kappa}$, the so-called *in-sample performance* of (8) is $\min_{v \in \mathcal{V}} \mathbb{E}_{\hat{\mathbb{P}}_\kappa}[f(\kappa, v)]$, and admits the following distributionally robust, semi-infinite variation over the Wasserstein ball centred at the empirical distribution $\hat{\mathbb{P}}_\kappa$:

$$\inf_{v \in \mathcal{V}} \sup_{\mathbb{Q} \in \mathcal{B}_\varepsilon(\hat{\mathbb{P}}_\kappa)} \mathbb{E}_{\mathbb{Q}}[f(\kappa, v)]. \quad (9)$$

The optimal value of (9) is known to upper bound the *out-of-sample performance*, $\mathbb{E}_{\mathbb{P}_\kappa}[f(\kappa, v)]$, with high confidence [6], [19]. Note that $\mathbb{E}_{\mathbb{P}_\kappa}[f(\kappa, v)]$ denotes the quantity of interest in studying (5), as it depends on the unknown distribution \mathbb{P}_κ . We will assume that this distribution is light-tailed, which is key to the results in [6], [19]:

Assumption 1: There exists some $a > 0$ such that $\mathbb{E}_{\mathbb{P}_\kappa}[e^{\|\kappa\|^a}] := \int_{\Theta} e^{\|\kappa\|^a} \mathbb{P}_\kappa(d\kappa) < \infty$. \square

Under Assumption 1, which is satisfied automatically if Θ is compact, [19, Th. 3.5] establishes that for any given confidence parameter $\beta > 0$, there exists some data-driven ambiguity radius, $\varepsilon = \varepsilon(\beta) > 0$, which guarantees the following probabilistic bound

$$\mathbb{P}_{\kappa}^{K_i+K} \left\{ \mathbb{E}_{\mathbb{P}_{\kappa}}[f(\kappa, v)] \leq \sup_{\mathbb{Q} \in \mathcal{B}_{\varepsilon}(\hat{\mathbb{P}}_{\kappa})} \mathbb{E}_{\mathbb{Q}}[f(\kappa, v)] \right\} \geq 1 - \beta. \quad (10)$$

However, solving the problem (9) is not trivial to solve since it is semi-infinite. Using the results in [19], [6, Th. 4.2] shows that considering $\mathcal{B}_{\varepsilon}(\hat{\mathbb{P}}_{\kappa})$ as an ambiguity set in (9) allows for a finite, convex reformulation. Specifically, the optimal value of (9) is upper bounded by

$$\min_{v \in \mathcal{V}} \left[f(\hat{\kappa}, v) + \varepsilon \cdot \max \left(\sup_{\xi \in \Xi} \|\xi\|_{\infty} \|\text{col}(g, 0)\|_*, \rho \|v\|_* \right) \right]. \quad (11)$$

Thus, by denoting v^* as an optimal solution to (11), the control law $u^* = \hat{\mathcal{U}}_f g^*$ enjoys the data-driven probabilistic guarantees in (10), obtaining good control performance with respect to the possible realizations of the stochastic output trajectory y associated with the ambiguity set $\mathcal{B}_{\varepsilon}(\hat{\mathbb{P}}_{\kappa})$.

Some consideration of the optimization problem in (11) is in order. First, we note that the cost function is convex since it corresponds to the sum of a (separable) convex function, f , and the pointwise maximum between dual norms, and is therefore also convex. Moreover, by defining $\zeta(\cdot) := \rho \|\cdot\|_1$, $\max(\sup_{\xi \in \Xi} \|\xi\|_{\infty} \|\text{col}(g, 0)\|_*, \rho \|v\|_*)$ is equivalent to

$$\begin{cases} \inf_{\lambda \geq 0} & \lambda \varepsilon \\ \text{s.t.} & \sup_{\xi \in \Xi} \|\xi\|_{\infty} \|\text{col}(g, 0)\|_* \leq \lambda, \\ & \sup_{\xi \in \Xi'} \|\xi\|_{\infty} \|v\|_* \leq \lambda, \end{cases}$$

where $\Xi' := \{\xi \in \mathbb{R}^{\ell K} \mid \zeta^*(\xi) < \infty\}$ is a bounded set. In fact, by the definition of the conjugate function, $\zeta^*(\xi) = 0$ if $\|\xi\|_{\infty} \leq \rho$, while $\zeta^*(\xi) = \infty$ otherwise. It therefore follows that (11) amounts to solving a conic optimization problem where the dimension of the decision variable v is $N - (K_i + K) + 1$, and N is a design parameter chosen so that $N \geq (m + 1)(n + K_i + K) - 1$, i.e., the persistent excitation assumption is satisfied. The large amount of data that one should collect may pose several challenges in the online implementation of distributionally robust controllers. Alternatively, one may also choose short control horizons K , resulting in poor control performance or instability of the controlled system. This motivates us to develop a procedure to compress the information brought by the dataset \mathcal{H}_{K_i+K} into a smaller synthetic set of representative system behaviours. In the next section, we formalize the data compression problem as a variational Wasserstein problem, proposing a solution procedure to design such a synthetic dataset.

IV. DATA COMPRESSION AS A VARIATIONAL WASSERSTEIN PROBLEM

In this section we propose an offline, optimal transport-based procedure that can compress a possibly large dataset of system trajectories to a smaller, synthetic one of representative behaviours, i.e., to select a limited set of synthetic representative samples that “best” summarize the information content. This clearly affects the control problem addressed by imposing a lower computational burden in solving an optimization problem similar to (11), hence making distributionally robust control approaches more appealing for online implementation. Moreover, we show that the optimal solution obtained by means of our synthetic dataset enjoys probabilistic guarantees of the same type in (10) on a Wasserstein ball with an enlarged ambiguity radius. Essentially, the marginal increase in the radius of the Wasserstein ball depends on the number of samples adopted, and vanishes as the size of the synthetic dataset grows to that of the original dataset.




Fig. 1: Schematic two-dimensional representation of the Wasserstein barycentres problem in (13). Each synthetic atom (red star) identifies a specific barycentre associated to a cluster (dashed red lines) of original samples (blue circles).

Let the dimensions of a matrix of input/output measurements over some horizon $L \in \mathbb{N}$, $\mathcal{H}_L \in \mathbb{R}^{r \times R}$, be fixed, i.e., $N \geq (m+1)(n+L) - 1$ be chosen so that the control input \hat{u} is persistently exciting of order $n+L$ as in Lemma 1, where $r := (m+\ell)L$, $R := N - L + 1$. The empirical distribution of such measurements is defined as $\hat{\mathbb{P}}_\kappa = \frac{1}{R} \sum_{i \in \mathcal{R}} \delta_{h_i}$, where $\mathcal{R} := \{1, \dots, R\}$ and h_i is the i -th column of \mathcal{H}_L . Our goal is to find a set of locations $\mathcal{S}_L := [s_1 \dots s_S] \in \mathbb{R}^{p \times S}$, with $S \leq R$, whose empirical probability distribution $\hat{\mathbb{P}}_s = \frac{1}{S} \sum_{i \in \mathcal{S}} \delta_{s_i}$, $\mathcal{S} := \{1, \dots, S\}$, is closest to that of the original dataset. Hence, our problem can be formulated as an optimal transport problem

$$\min_{\mathcal{S}_L} \mathcal{W}(\hat{\mathbb{P}}_\kappa, \hat{\mathbb{P}}_s) = \min_{\mathcal{S}_L} \min_{T \in \mathcal{T}(\mathbf{1}_{R/R}, \mathbf{1}_{S/S})} \langle T, D(\mathcal{H}_L, \mathcal{S}_L) \rangle. \quad (12)$$

Our optimization problem has a strong practical interpretation. Specifically, an optimal solution to (12), \mathcal{S}_L^* , is one that produces a distribution $\hat{\mathbb{P}}_s$ closest to the original one in the Wasserstein distance, and is therefore the one that minimizes the transport cost between the two distributions $\hat{\mathbb{P}}_s$ and $\hat{\mathbb{P}}_\kappa$. With a slight abuse of notation, we define by $\eta(S)$ the optimal value to (12), which clearly depends on S , the number of samples adopted. This quantity is key in characterizing the robustness properties of the control approach that we are about to introduce. The nested optimization program in (12) is a variational Wasserstein problem, representing a particular case of the Wasserstein barycentres problem [20]. Specifically, it directly falls into the set of k -means problems [21]. In addition, for semi-discrete settings the benefit of adopting the Wasserstein distance as a metric to compare probability measures has been proved in many theoretical and practical problems, from dictionary and statistical learning [22], [23], to vision and image processing [24], [25].

A. On the optimal transport problem (12)

Despite its appealing structure and strong practical interpretation, the variational Wasserstein problem in (12) is convex in each single variable, i.e., \mathcal{S}_L and T , but not jointly. In practice, this might produce locally optimal solutions, where every atom defining each optimal \mathcal{S}_L^* identifies a specific barycentre

for a subset of samples in \mathcal{H}_L . Specifically, in our setting, (12) corresponds to

$$\min_{\mathcal{S}_L} \min_{T \in \mathcal{T}(\mathbf{1}_R/R, \mathbf{1}_S/S)} \sum_{j \in \mathcal{S}} \sum_{i \in \mathcal{R}} t_{i,j} \|h_i - s_j\|. \quad (13)$$

For any fixed $j \in \mathcal{S}$, every $t_{i,j} \geq 0$ defines the quantity of a predefined sample h_i of \mathcal{H}_L that is associated with a barycentre s_j – see, e.g., Fig.1 for a schematic representation. In this specific example, each sample h_i is associated to the j -th barycentre only, i.e., $t_{i,j} = 1$ and $t_{i,h} = 0$, for all $h \in \mathcal{S} \setminus \{j\}$, meaning that there are no overlapping clusters. In principle, as $t_{i,j} \geq 0$, the information associated with any sample h_i may be distributed among several barycentres, thereby producing overlapping clusters. It is easily shown that every atom defining \mathcal{S}_L^* belongs to $\text{conv}(\mathcal{H}_L)$.

Since the inner minimization problem in (13) represents the pointwise minimum of linear functions, we note that the Wasserstein distance is not smooth in its arguments. To circumvent this problem, the cost function in (12) can be regularized by means of a strictly convex, weighted entropic term, i.e., $\gamma \langle T, \log(T) \rangle$, for some $\gamma > 0$. The benefits are twofold [26]: i) the inner optimization problem in (12) admits a closed form, which translates to a matrix balancing problem, and ii) the Wasserstein distance is differentiable. To solve (13) one may also rely on the convexity in each single variable of the Wasserstein distance, for which possible solution algorithms are, e.g., the typical alternate block-coordinate descent methods [27], [28].

B. Robust performance guarantees

From [19, Th. 3.5], we know that, for any confidence parameter β , there exists some $\varepsilon > 0$, that depends only on the amount of available data, such that the optimal value of (9) upper bounds the out-of-sample $\mathbb{E}_{\mathbb{P}_\kappa}[f(\kappa, v)]$. Our approach then amounts to first solving the optimal transport problem in (12), computing a synthetic set of atoms \mathcal{S}_L , and then reformulating the robust optimization problem in (9) with a Wasserstein ball centred on $\hat{\mathbb{P}}_s$ instead of $\hat{\mathbb{P}}_\kappa$, i.e.,

$$\inf_{v \in \mathcal{V}} \sup_{\mathbb{Q} \in \mathcal{B}_\varepsilon(\hat{\mathbb{P}}_s)} \mathbb{E}_{\mathbb{Q}}[f(\kappa, v)]. \quad (14)$$

The fact that $S < R$ in general, namely we are designing a reduced set of synthetic samples relative to the original dataset, intuitively has two main implications:

- i) The optimal value in (14) can still achieve the performance bound in (10) with high confidence, at a price of considering a larger radius of the ambiguity set $\mathcal{B}_\varepsilon(\hat{\mathbb{P}}_s)$ (see Proposition 1). We show that the additional term accounting for additional robustness vanishes as the number of atoms tends to the cardinality of the original dataset;
- ii) The robust optimization problem in (14) can be manipulated to obtain a tractable convex reformulation equivalent to the one in (11), but defined on a lower dimensional space. In fact, while the optimization variable in (11) has dimension $R = N - L + 1$, where N is chosen so that the condition on the persistency of excitation is met, which leads to $R \geq (n+1) + \frac{1}{m+1}(mN - 1)$, the finite, convex formulation obtained by manipulating (14) establishes that v simply belongs to \mathbb{R}^S .

We remark that, instead, the number of constraints defining \mathcal{V} does not change since it depends on the control horizon L (design parameter). Despite this dimensionality mismatch, for simplicity's sake we will keep the same notation in the rest of the paper.

Next, by relying on the definition of $\eta(S)$ following (12), we show that an optimal solution to the robust optimization problem in (14) upper bounds the out-of-sample performance $\mathbb{E}_{\mathbb{P}_\kappa}[f(\kappa, v)]$ with high confidence.

Proposition 1: Let $\beta \in (0, 1)$ be some given confidence parameter, and let $S \leq R$ be fixed. Under Assumption 1, there exists some $\bar{\varepsilon} = \bar{\varepsilon}(\beta, S) > 0$ such that, for all $v \in \mathcal{V}$,

$$\mathbb{P}_\kappa^S \left\{ \mathbb{E}_{\mathbb{P}_\kappa}[f(\kappa, v)] \leq \sup_{\mathbb{Q} \in \mathcal{B}_{\bar{\varepsilon}}(\hat{\mathbb{P}}_s)} \mathbb{E}_{\mathbb{Q}}[f(\kappa, v)] \right\} \geq 1 - \beta.$$

Algorithm 1: Receding horizon robust synDeePC

Offline: Given \mathcal{H}_{K_i+K} , set $S \leq N - (K_i + K) + 1$, compute $\mathcal{S}_{K_i+K}^* \in \operatorname{argmin}_{\mathcal{S}_{K_i+K}} \mathcal{W}(\hat{\mathbb{P}}_\kappa, \hat{\mathbb{P}}_s)$

Initialization: Set $\hat{\mathcal{V}}(0)$, $\hat{\kappa}(0)$ and $\bar{\varepsilon}$

Iteration ($k \in \mathbb{N}$):

(S1) Compute

$$v^*(k) := \operatorname{argmin}_{v \in \mathcal{V}(k)} [f(\hat{\kappa}(k), v) + \bar{\varepsilon} \cdot \max(\sup_{\xi \in \Xi} \|\xi\|_\infty \|\operatorname{col}(g, 0)\|_*, \rho \|v\|_*)]$$

(S2) Set $u^*(k) = \mathcal{U}_f v^*(k)$, apply $u_1^*(k)$

(S3) Collect measurements, update $\mathcal{V}(k+1)$, $\hat{\kappa}(k+1)$

TABLE I: Main simulation parameters

K	K_i	N	T_s	c	ρ	$\epsilon(\beta)$
30	1	214	0.05	200	10^5	10^{-3}

□

Proof: First, given any $S \leq R$, we recall that $\eta(S)$ corresponds to the Wasserstein distance between the discrete probability distribution associated with the original dataset, $\hat{\mathbb{P}}_\kappa$, and the computed synthetic one, $\hat{\mathbb{P}}_s$. Thus, the triangle inequality for the Wasserstein metric ensures that the distance between the real, unknown distribution \mathbb{P}_κ and $\hat{\mathbb{P}}_s$ can be upper bounded as follows

$$\mathcal{W}(\mathbb{P}_\kappa, \hat{\mathbb{P}}_s) \leq \mathcal{W}(\mathbb{P}_\kappa, \hat{\mathbb{P}}_\kappa) + \mathcal{W}(\hat{\mathbb{P}}_\kappa, \hat{\mathbb{P}}_s) = \mathcal{W}(\mathbb{P}_\kappa, \hat{\mathbb{P}}_\kappa) + \eta(S).$$

Moreover, in view of Assumption 1, it follows from [19, Th. 3.4] that, for any fixed $\beta \in (0, 1)$, there exists some $\varepsilon(\beta) > 0$ such that $\mathbb{P}_\kappa^R\{\mathcal{W}(\mathbb{P}_\kappa, \hat{\mathbb{P}}_\kappa) \leq \varepsilon(\beta)\} \geq 1 - \beta$. Therefore, since $\hat{\mathbb{P}}_s$ is an empirical distribution as well, we obtain $\mathbb{P}_\kappa^S\{\mathcal{W}(\mathbb{P}_\kappa, \hat{\mathbb{P}}_s) \leq \varepsilon(\beta) + \eta(S)\} \geq 1 - \beta$. This latter relation, which can be equivalently restated as $\mathbb{P}_\kappa^S\{\mathbb{P}_\kappa \in \mathcal{B}_{\bar{\varepsilon}}(\hat{\mathbb{P}}_s)\} \geq 1 - \beta$, where $\bar{\varepsilon} := \varepsilon(\beta) + \eta(S)$, directly implies $\mathbb{E}_{\mathbb{P}_\kappa}[f(\kappa, v)] \leq \sup_{\mathbb{Q} \in \mathcal{B}_{\bar{\varepsilon}}(\hat{\mathbb{P}}_s)} \mathbb{E}_{\mathbb{Q}}[f(\kappa, v)]$ with probability $1 - \beta$, thus concluding the proof. ■

We note that the Wasserstein distance between $\hat{\mathbb{P}}_\kappa$ and the empirical distribution associated with the compressed dataset, $\hat{\mathbb{P}}_s$, can be made arbitrarily small by increasing the number of atoms S , since $\min_{\mathcal{S}_L} \mathcal{W}(\hat{\mathbb{P}}_\kappa, \hat{\mathbb{P}}_s) \rightarrow 0$ as $S \rightarrow R$, and hence $\eta(S) \rightarrow 0$. In this case, we recover the radius of the ambiguity set in [6, Th. 4.1], although the behaviour of $\eta(S)$ is not monotonically decreasing to 0 as we will see in the next section. Finally, we remark that the optimization problem in (14) admits a tractable reformulation identical to the one in (11), but which can be solved online with significantly lower computational burden. This then paves the way to adopt possibly longer control horizons L , thus enhancing the control performance and without compromising the closed-loop stability of the system. These aspects are investigated in the next section.

V. NUMERICAL SIMULATIONS

In this section we apply the data-enabled predictive control (DeePC) method presented in [6] with an additional, offline step that computes a synthetic dataset, thus renamed synDeePC (Algorithm 1). Specifically, we compare control and computational performance when considering the original dataset and a compressed one when steering a linear model of a quadcopter in a receding horizon fashion.




Fig. 2: Wasserstein distance vs. number of synthetic atoms.

Simulations are run in Matlab by using Gurobi [29] as a solver for (S1) in Algorithm 1, on a laptop with a Quad-Core Intel Core i5 2.4 GHz CPU and 8 Gb RAM. The main parameters adopted are summarized in Table I.

The linear model adopted is valid around a hover position, where the state vector is $\text{col}(x, y, z, \dot{x}, \dot{y}, \dot{z}, \phi, \theta, \psi, \dot{\phi}, \dot{\theta}, \dot{\psi}) \in \mathbb{R}^{12}$. Here, x , y and z are the three spatial coordinates and relative velocities $(\dot{x}, \dot{y}, \dot{z})$, while ϕ , θ , and ψ are the angular ones, with relative rates $(\dot{\phi}, \dot{\theta}, \dot{\psi})$. The control inputs are represented by four identical rotors, constrained to the set $\mathcal{U} = [-0.7007, 0.2993]^{4(K_i+K)}$ due to physical limitations. By assuming full state measurement, we use the same state-space matrices adopted in [6, §V], as well as same original cost function, $J(u, y) = \|u\|_1 + c\|y - r\|_1$, where r denotes a parametrized, 8-figure trajectory with fixed altitude. The parameter T_s in Table I represents the temporal resolution with which the reference trajectory is sampled (i.e., the sampling time). Moreover, with the adopted values, the 184 columns of the matrix \mathcal{H}_{31} are filled by means of random inputs drawn from a uniform distribution on \mathcal{U} : this is to guarantee the persistency of excitation for the (syn)DeePC, according to Definition 2.

Some a-priori considerations on the choice of the parameter S in Algorithm 1 can be made, e.g., in a data-driven fashion. Specifically, we evaluate the behaviour of the Wasserstein distance $\mathcal{W}(\hat{\mathbb{P}}_\kappa, \hat{\mathbb{P}}_s)$ when S varies, solving the offline step without regularization by means of a standard block-descent algorithm [27]. Thus, according to Fig. 2, the function $\eta(S)$ takes reasonable values for $S \leq 92$, leading to an offline step in Algorithm 1 taking less than three minutes. Interestingly, we note that the effect of the local minima seems to prevent the Wasserstein distance from decreasing monotonically for values of $S > 92$, i.e., $R/2$, as in Fig. 2.

In Fig. 3 and 4, we compare the trajectory tracking performances of the quadrotor controlled by means of the DeePC with full dataset \mathcal{H}_{31} (solid lines) and synthetic dataset, \mathcal{S}_{31} , obtained first by reducing to the 50% the total number of samples, i.e., $S = 92$ (solid-dashed lines), which also correspond to a reduction of the 70% w.r.t. the total number of samples when considering a longer control horizon, i.e., $K = 50$ (dotted lines). An example of a typical constrained input signal for the rotors can be found in Fig. 5, where the behaviour of the first element of u^* in all three cases is shown. Here, we used an equivalent statistic for the noise acting on each measurement channel, i.e., $\nu \sim \mathcal{N}(0, 2^{-7})$, whose value is chosen to match the experimental setup in [30]. As shown in Fig. 4, where the position errors of the spatial coordinates are illustrated, the performances of the robust controller computed by means of the synthetic dataset \mathcal{S}_{31} do not degrade markedly compared with the one computed by means of \mathcal{H}_{31} , also exhibiting an almost overlapping behaviour when the control horizon K increases (\mathcal{S}_{51}). Moreover, from our numerical experience, the step (S1) in Algorithm 1 with a compressed dataset takes approximately




Fig. 3: Dynamical evolution of the controlled quadcopter while following a figure-8 trajectory.




Fig. 4: Spatial coordinate tracking errors.




Fig. 5: First element of the control input u^* .




Fig. 6: Computational time with Gurobi [29] over the whole trajectory tracking control problem.




Fig. 7: Trajectory tracking performance of the controlled quadcopter with a higher level of noise, $\nu \sim \mathcal{N}(0, 1^{-3})$.

0.64[s] on average, in sharp contrast to the 1.73[s] required by the original dataset, see Fig. 6. On the other hand, a longer control horizon does not lead to a much higher computational time, i.e., 0.85[s], while considering the whole dataset would take around 3[s] to solve the DeePC optimization problem.

Finally, we investigate how a higher level of noise acting on the measurements reflects on the control performance. Specifically, we assume an equivalent statistic on the channels, i.e., $\nu \sim \mathcal{N}(0, 1^{-6})$. In this case, with $N = 214$ and main parameters as in Tab. I, we experienced that both controllers do not accomplish the trajectory tracking problem, directly leading to instability. For this reason, we decide to collect $3N$ measurements. As shown in Fig. 7, while the original DeePC is not able to follow the figure-8 reference from the very beginning (solid lines), the synDeePC in Algorithm 1 with $S = 306$ introduces an offset on the y -axis tracking only, while keeping good performance on both x and z axes (solid-dashed lines).

VI. CONCLUSION AND OUTLOOK

The optimal transport approach promises to be a key tool for the design of synthetic datasets guaranteeing both robust performance and reasonable computational burden for real-time implementation

of data-driven controllers. Specifically, we have investigated the benefits of adopting the Wasserstein metric to compress the informative content of a large dataset into a smaller one, also illustrating the performance of the robust controller obtained by means of the synthetic dataset compared to the original one. Future research directions will focus on the impact that the discrete measure adopted to compare the empirical distribution associated with the original data, i.e., the vector β in (2), has on the robustness. Intuitively, if one was to allow variable weights in β , we envision that Fig. 2 would be necessarily monotone, assuming one could solve to global optimality. The reason is that the $(S + 1)$ -th point could always be placed at an arbitrary location with zero mass, and this would require the same Wasserstein distance as with S synthetic points only. It would also mean that $\eta(S) = 0$ for all $S > R$. Moreover, given the input/output structure of the gathered data, we will investigate also the possibility to use different distances to define the Wasserstein metric, as well as a jointly convex reformulation of the whole program in (12).

REFERENCES

- [1] J. C. Willems, P. Rapisarda, I. Markovsky, and B. L. De Moor, "A note on persistency of excitation," *Systems & Control Letters*, vol. 54, no. 4, pp. 325–329, 2005.
- [2] C. De Persis and P. Tesi, "Formulas for data-driven control: Stabilization, optimality, and robustness," *IEEE Transactions on Automatic Control*, vol. 65, no. 3, pp. 909–924, 2019.
- [3] H. J. Van Waarde, J. Eising, H. L. Trentelman, and M. K. Camlibel, "Data informativity: a new perspective on data-driven analysis and control," *IEEE Transactions on Automatic Control*, 2020.
- [4] J. Berberich, A. Koch, C. W. Scherer, and F. Allgöwer, "Robust data-driven state-feedback design," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 1532–1538.
- [5] J. Coulson, J. Lygeros, and F. Dörfler, "Data-enabled predictive control: In the shallows of the DeePC," in *2019 18th European Control Conference (ECC)*. IEEE, 2019, pp. 307–312.
- [6] J. Coulson, J. Lygeros, and F. Drfler, "Regularized and distributionally robust data-enabled predictive control," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2019, pp. 2696–2701.
- [7] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Data-driven model predictive control with stability and robustness guarantees," *IEEE Transactions on Automatic Control*, 2020.
- [8] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, 2000.
- [9] P. L. Bartlett, M. I. Jordan, and J. D. McAuliffe, "Convexity, classification, and risk bounds," *Journal of the American Statistical Association*, vol. 101, no. 473, pp. 138–156, 2006.
- [10] J. Wang and S. J. Qin, "A new subspace identification approach based on principal component analysis," *Journal of process control*, vol. 12, no. 8, pp. 841–855, 2002.
- [11] —, "Closed-loop subspace identification using the parity space," *Automatica*, vol. 42, no. 2, pp. 315–320, 2006.
- [12] B. Schölkopf, A. Smola, and K.-R. Müller, "Kernel principal component analysis," in *International conference on artificial neural networks*. Springer, 1997, pp. 583–588.
- [13] P. Van Overschee and B. De Moor, "Closed loop subspace system identification," in *Proceedings of the 36th IEEE Conference on Decision and Control*, vol. 2. IEEE, 1997, pp. 1848–1853.
- [14] T. McKelvey, H. Akçay, and L. Ljung, "Subspace-based multivariable system identification from frequency response data," *IEEE Transactions on Automatic Control*, vol. 41, no. 7, pp. 960–979, 1996.
- [15] M. Jansson and B. Wahlberg, "On consistency of subspace methods for system identification," *Automatica*, vol. 34, no. 12, pp. 1507–1519, 1998.
- [16] C. Villani, *Topics in optimal transportation*. American Mathematical Society, 2003, no. 58.
- [17] G. Peyré and M. Cuturi, "Computational optimal transport," *Foundations and Trends® in Machine Learning*, vol. 11, no. 5-6, 2019.
- [18] A. Mesbah, "Stochastic model predictive control: An overview and perspectives for future research," *IEEE Control Systems Magazine*, vol. 36, no. 6, pp. 30–44, 2016.
- [19] P. M. Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations," *Mathematical Programming*, vol. 171, no. 1-2, pp. 115–166, 2018.
- [20] M. Cuturi and A. Doucet, "Fast computation of Wasserstein barycenters," *Journal of Machine Learning Research*, 2014.
- [21] M. K. Ng, "A note on constrained k-means algorithms," *Pattern Recognition*, vol. 33, no. 3, pp. 515–519, 2000.
- [22] A. Rolet, M. Cuturi, and G. Peyré, "Fast dictionary learning with a smoothed Wasserstein loss," in *Artificial Intelligence and Statistics*, 2016, pp. 630–638.
- [23] C. Frogner, C. Zhang, H. Mobahi, M. Araya, and T. A. Poggio, "Learning with a Wasserstein loss," in *Advances in Neural Information Processing Systems*, 2015, pp. 2053–2061.
- [24] M. Cuturi and G. Peyré, "A smoothed dual approach for variational Wasserstein problems," *SIAM Journal on Imaging Sciences*, vol. 9, no. 1, pp. 320–343, 2016.
- [25] J. Lellmann, D. A. Lorenz, C. Schonlieb, and T. Valkonen, "Imaging with Kantorovich–Rubinstein discrepancy," *SIAM Journal on Imaging Sciences*, vol. 7, no. 4, pp. 2833–2859, 2014.

- [26] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," in *Advances in neural information processing systems*, 2013, pp. 2292–2300.
- [27] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: numerical methods*. Prentice Hall Englewood Cliffs, NJ, 1989, vol. 23.
- [28] A. Beck and L. Tretuashvili, "On the convergence of block coordinate descent type methods," *SIAM Journal on Optimization*, vol. 23, no. 4, pp. 2037–2060, 2013.
- [29] L. Gurobi Optimization, "Gurobi optimizer reference manual," 2020. [Online]. Available: <http://www.gurobi.com>
- [30] E. Elokda, J. Coulson, P. Beuchat, J. Lygeros, and F. Dörfler, "Data-enabled predictive control for quadcopters," 2019, ETH Zurich, Automatic Control Laboratory.