Non-Parametric and Regularized Dynamical Wasserstein Barycenters for Time-Series Analysis

Kevin C. Cheng*, Shuchin Aeron*, Michael C. Hughes†, Eric L. Miller*

arXiv:2210.01918v2 [cs.LG] 7 Oct 2022

Abstract

We consider probabilistic time-series models for systems that gradually transition among a finite number of states. We are particularly motivated by applications such as human activity analysis where the observed time-series contains segments representing distinct activities such as running or walking as well as segments characterized by continuous transition among these states. Accordingly, the dynamical Wasserstein barycenter (DWB) model introduced in Cheng et al. in 2021 [1] associates with each state, which we call a *pure state*, its own probability distribution, and models these continuous transitions with the dynamics of the barycentric weights that combine the pure state distributions via the Wasserstein barycenter. Here, focusing on the univariate case where Wasserstein distances and barycenters can be computed in closed form, we extend [1] by discussing two challenges associated with learning a DWB model and two improvements. First, we highlight the issue of uniqueness in identifying the model parameters. Secondly, we discuss the challenge of estimating a dynamically evolving distribution given a limited number of samples. The uncertainty associated with this estimation may cause a model's learned dynamics to not reflect the gradual transitions characteristic of the system. The first improvement introduces a regularization framework that addresses this uncertainty by imposing temporal smoothness on the dynamics of the barycentric weights while leveraging the understanding of the non-uniqueness of the problem. This is done without defining an entire stochastic model for the dynamics of the system as in [1]. Our second improvement lifts the Gaussian assumption on the pure states distributions in [1] by proposing a quantile-based non-parametric representation. We pose model estimation in a variational framework and propose a finite dimensional approximation to the infinite dimensional problem, which is then solved using constrained coordinate descent. The discussed challenges and improvements are demonstrated through the ability of our non-parametric, regularized DWB model to accurately characterize the distribution of the data as the system transition among its dynamical modes for simulated and real world human activity time-series.

Index Terms

Wasserstein barycenter, displacement interpolation, dynamical model, time-series analysis, sliding window, non-parametric, quantile function, human activity analysis.

I. Introduction

We consider a probabilistic model for time-series data where the observation at each point in time depends on a dynamically evolving latent state. We are particularly motivated by systems that continuously move among a set of canonical behaviors, which we call *pure states*. Over some periods, the system may in reside entirely in one of the pure state while over other periods, the system is transitioning among these pure states in a temporally smooth manner. There are many applications where such a model is appropriate including climate modeling [2], sleep analysis [3], simulating physical systems [4], as well as characterizing human activity from video [5] or wearable-derived accelerometry [6] data. Using the last case as an example, over some periods of time the individual may be engaged in a well-defined activity such as standing or running, during which the data can be taken to be distributed according to a probability distribution corresponding to that pure state. Given the high sampling rates of modern sensors, there also may be time intervals where multiple consecutive observations reflect the gradual transition between or among pure states, during which the distribution of the data is given by a suitable combination of the pure state distributions. Therefore, one possible model for these types of systems consists of three components: a set of distributions that corresponds to the data while the system resides in each pure state, a dynamically evolving latent state which captures the transition dynamics of the system as it moves among these pure states, and a means of interpolating among these pure state distributions that characterizes the data distribution in the transition regions.

These types of systems pose some unique considerations that are not sufficiently addressed by prior work in time-series modeling. The two most common methods for modeling latent state systems are continuous and discrete state-space models. Continuous state-space models [7], [8], [9] have no natural way to identify these pure states in which the system may persist for periods of time. Doing so would require some form of clustering or other analysis on the latent state-space. In discrete state-space models such as hidden Markov models, [10], [11], [12], the latent state vector represents the probability that the system resides in each of a countable number of discrete (or in our terminology, pure) states. That is, the system is assumed to be in exactly one pure state at each time, but there is some uncertainty as to which one. For these models, the data-generating distribution given by the latent state vector is a convex combination of the pure state distributions. We argue in [1] this is
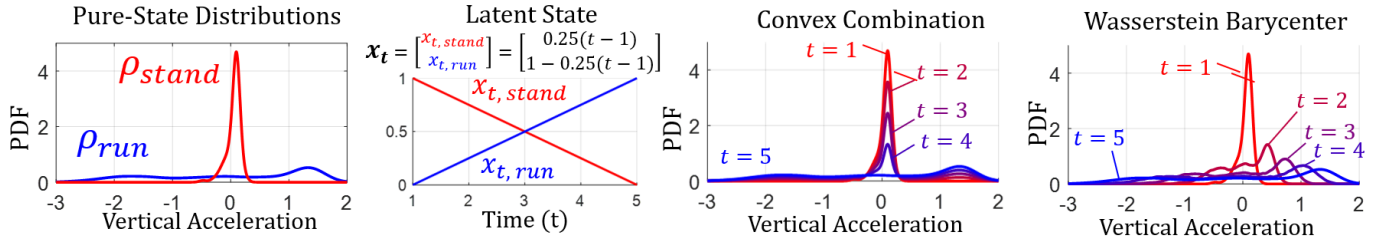
Fig. 1: **Comparison of convex combination vs Wasserstein barycenter for modeling human activity transitions**. The beep test (BT) dataset consists of a subject running back-and-forth between two points, stopping at each point (see Sec. V-B for more details). (a) The PDFs of the vertical acceleration the system's two pure states (stand, run) are estimated using a KDE with a Gaussian kernel whose mean corresponds to the observed data when the system resides in these pure states. Modeling a transition from stand to run via the time-varying weights (b) for $t = 1, .., 5$, we show the resulting data distributions during this transition region according to a convex combination (c) and Wasserstein barycenter (d) interpolation model.

also an insufficient model for the problems which interest us. Returning to our human activity example, the data distribution illustrated in Fig. 1 produced by a convex combination of the underlying "standing" and "running" pure state distributions can be interpreted as "sometimes standing" and "sometimes running," which is not a proper description of the gradual acceleration that actually occurs.

A better model for interpolating the data distribution during the transition period between standing and running would smoothly shift probability mass between the two pure state distributions. Again, shown by Fig. 1, such smooth blending can be achieved through displacement interpolation [13] via the *Wasserstein barycenter* [14]. Our recent work proposed a new time-series model, Dynamical Wasserstein Barycenters (DWB) [1], which models the dynamically evolving data distribution of time series as a sequence of Wasserstein barycenter distributions, parameterized by a set of pure state distributions and a sequence of barycentric weights. These weights which lie on the probability simplex are taken to be the latent state of the model. We fit a model by estimating the distribution of an observed time series over time from a sliding window of samples and minimizing the Wasserstein distance between the model specified Wasserstein barycentric distribution to this sequence of estimated distributions.

In this paper, we expand on [1] by highlighting two challenging characteristics of the DWB modeling problem and two improvements that address certain limitations of the original DWB approach. The first characteristic relates to uniqueness, where multiple combinations of pure state distributions and barycentric weights can produce the same Wasserstein barycentric distribution. Although this is true for multidimensional distributions, we use a univariate formulation to transparently demonstrate how this non-uniqueness is captured in an *inverse-scaling* relationship between the model's latent state and pure state parameters. For a distribution that is the Wasserstein barycenter of a set of pure states distributions and barycentric weights, we can construct the same barycenter distribution by suitably reducing (resp. increasing) the distances among the latent states on the simplex and correspondingly increasing (resp. decreasing) the Wasserstein distances among the pure state distributions.

For the second characteristic we explore the challenge of estimating a dynamically evolving data distribution from a single instance of a time series using a window of samples. The uncertainty of these window estimates may cause variability in learned latent state resulting in it not reflecting the temporal smoothness associated the system. We show how the accuracy and variability of such estimates depends on the size of the sample window and the magnitude of the change in distribution of the samples in the window. Regarding window size, smaller windows lack the samples for precise estimation of the data distribution at a current point in time. On the other hand, larger windows span a longer time period over which the data distribution may change, which reduces the accuracy of the estimated distribution. In a simulated example, where the dynamics consists of constant rate transitions between two Gaussian states, we show that there exists an optimal window size that balances these two effects and discuss the dependency of this window size tradeoff on the temporal dynamics and pure states of the system.

Our first improvement addresses the limitations of the choice of [1] to use a stochastic generative model to impose temporal smoothness on the system's latent state to mitigate the uncertainty associated with the window estimates. Such an approach may introduce additional unnecessary or potentially undesirable probabilistic properties on the latent state process such as a limiting distribution [15]. Instead, we propose a regularization scheme imposes temporal smoothness by directly penalizing the difference between the latent state vectors at adjacent point in time. Drawing from the field of compositional analysis [16], we employ the Bhattacharyya-arccos distances [17] as a distance measure well-suited to the probability simplex. As specified by the inverse-scaling relationship, introducing this latent state regularizer impacts the model's pure state distribution in a manner that causes them to diverge from the data. Therefore, we also introduce a regularizer to counteract this effect to ensure that the learned pure state distributions are representative of data while the system resides in each pure state.

Our second improvement addresses the Gaussian parametric limitation of [1] by proposing a non-parametric model for the pure state distributions. This extends the DWB framework to a wider class of time-series where the data distributions may not be Gaussian. We focus in this paper on the univariate case where in one-dimension the Wasserstein-2 distance between distributions is equivalent to the 2-norm between their respective quantile functions [18]. Using a discrete approximation to

the quantile function corresponding to the pure state distribution leads to a convenient linear least squares type model for the data misfit term in our estimation problem.

Our numerical experiments empirically validate our analysis and proposed improvements for the DWB model. Using simulated data, we demonstrate in a controlled setting how we effectively regularize our model parameters with proper consideration of the inverse-scaling analysis and the impact of window size on the accuracy of the model parameters. Additionally, using real world human activity data, we show how our non-parametric approach leads to improved estimation of the system's pure state distributions as well as improved fit of the time-evolving distribution of the observed data.

In summary, the primary contributions of this work consist of a discussion of two key challenging characteristics of the DWB modeling problem and two improvements that address the limitations of prior work:

1) We highlight the non-uniqueness of the parameters corresponding to a Wasserstein barycenter by detailing the inverse-scaling relationship between the pure state distributions and the simplex-valued barycentric weights.
2) We explore the impact of the window size on the ability to accurately estimate a dynamically evolving data distribution through exploring the tradeoff between the errors associated with large and small windows and the dependency of this tradeoff on the dynamics and pure states of the system.
3) We propose a regularization framework for the model parameters that imposes temporal smoothness in the latent states while accounting for the effect on the pure state distributions through inverse-scaling effect.
4) We propose a flexible non-parametric representation for the univariate pure state distributions using a discrete approximation to the quantile function that results in a finite dimensional formulation for DWB learning.

Our paper proceeds as follows: in Sec. II, we provide an overview of the Wasserstein distance and barycenter focusing on the univariate case. In Sec. III, we discuss the DWB model, highlighting the non-uniqueness and inverse-scaling property of the Wasserstein barycenter in as well as the impact of the window size on the estimation of a dynamically evolving data distribution. In Sec. IV, we develop our variational framework for learning a DWB model, followed by a discussion of our regularization framework for the learned model parameters, the discretization of the pure state distributions required to obtain a finite dimensional estimation problem. We then formally state our non-parametric and regularized DWB variational learning problem and provide an algorithm to learn the model parameters. In Sec. V we use simulated data to demonstrate the non-uniqueness, impact of window size and regularization terms discussed in this work and use real world human activity data to demonstrate the advantages of the non-parametric DWB approach relative to the Gaussian model.

II. TECHNICAL BACKGROUND

The Wasserstein-2 distance is a metric on the space of probability distributions on $\mathbb{R}^d$ with finite second moment [18]. For two random variables q and s with distributions ρ_q and ρ_s respectively, the squared Wasserstein-2 distance with square-Euclidean ground cost is defined via,

$$\mathcal{W}_2^2(\rho_q, \rho_s) = \inf_{\pi \in \Pi(\rho_q, \rho_s)} \mathbb{E}_{q,s \sim \pi} \|q - s\|_2^2 \tag{1}$$

where π denotes the joint distribution of q and s, and $\Pi(\rho_q, \rho_s)$ is the set of all joint distributions with marginals ρ_q, ρ_s. In this work, we refer to Eq. (1) as the squared Wasserstein distance.

Given a set of distributions $\rho_{q_{1:K}} = \{\rho_{q_1}, \rho_{q_2}, ..., \rho_{q_K}\}$ and a vector $\boldsymbol{x} \in \Delta^K$, where Δ^K denotes the standard K-simplex, the *Wasserstein barycenter* is the distribution that minimizes the weighted (with respect to elements in $\boldsymbol{x}$) squared Wasserstein distance to the set of distributions [14] and is given by,

$$\rho_B = B(\boldsymbol{x}, \rho_{q_{1:K}}) = \operatorname*{argmin}_{\rho} \sum_{k=1}^{K} \boldsymbol{x}[k] \mathcal{W}_2^2(\rho, \rho_{q_k}). \tag{2}$$

This minimization occurs over all probability distributions with finite second moment [14].

When ρ_q and ρ_s are univariate distributions with cumulative distribution functions P_q, P_s, the squared Wasserstein distance in Eq. (1) becomes [19],

$$\mathcal{W}_2^2(\rho_q, \rho_s) = \int_0^1 \left(P_q^{-1}(\xi) - P_s^{-1}(\xi) \right)^2 d\xi. \tag{3}$$

Here P_q^{-1} and P_s^{-1} are quantile functions, the generalized inverse [20] of the distribution function, given by,

$$P^{-1}(\xi) = \inf\{g \in \mathbb{R} : P(g) \geq \xi\}. \tag{4}$$

It follows from Eq. (3) and Eq. (2) that the Wasserstein barycenter of a set of univariate distributions with quantile functions $P_{q_{1:K}}^{-1}$, will have quantile function [19],

$$P_B^{-1} = \sum_{k=1}^{K} \boldsymbol{x}[k] P_{q_k}^{-1}. \tag{5}$$

III. THE DYNAMICAL WASSERSTEIN BARYCENTER MODEL

Shown in Fig. 2, the DWB model [1] describes the distribution of a time-series at time t as,

$$\rho_{B_t} = B(\boldsymbol{x}_t, \rho_{q_{1:K}}) \tag{6}$$

where ρ_{q_k}, $k = 1, 2, \ldots, K$ are the distributions of the pure states and the barycentric weight $\boldsymbol{x}_t \in \Delta^K$ capture the dynamics of the transitions among these pure states. Learning the parameters of a DWB model requires estimating the pure state distributions
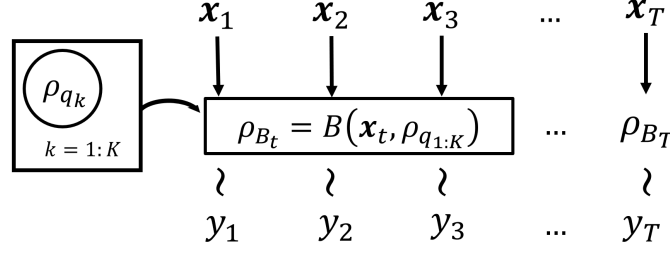
Fig. 2: **DWB model diagram**. The DWB models the distribution ρ_{B_t} from which the time-series y_t is sampled as the Wasserstein barycenter of a set of pure state distributions $\rho_{q_{1:K}}$ and barycentric weight $\boldsymbol{x}_{1:T}$, the latent state of the model.
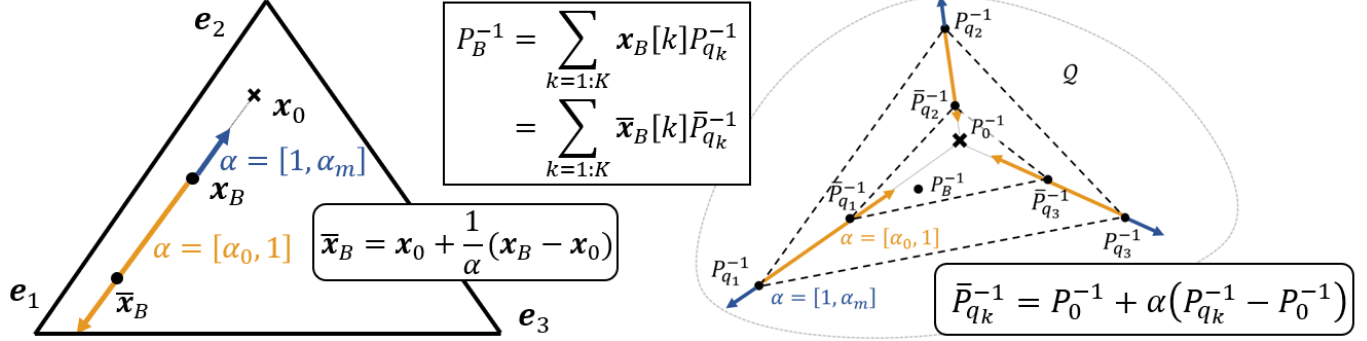

Fig. 3: **Diagram of the non-uniqueness and inverse-scaling effect of the parameters of a Wasserstein barycenter.** Consider a set of three pure states with quantile functions $P_{q_{1:3}}^{-1}$ and simplex-valued weight $\boldsymbol{x}_B \in \Delta^3$ where $\rho_B = B(\boldsymbol{x}_B, \rho_{q_{1:3}})$ with quanitle function $P_B^{-1} = \sum_{k=1}^3 \boldsymbol{x}_B[k] P_{q_k}^{-1}$. We construct a family of of distinctly different pure state quantile functions $\bar{P}_{q_{1:K}}^{-1}$ and barycentric weights $\bar{\boldsymbol{x}}_B$ which produce the exact same barycentric distribution $P_B^{-1} = \sum_{k=1}^3 \bar{\boldsymbol{x}}_B[k] \bar{P}_{q_k}^{-1}$. Let $\boldsymbol{x}_0$ be another point on the simplex where $\rho_0 = B(\boldsymbol{x}_0, \rho_{q_{1:3}})$ has quantile function $P_0^{-1} = \sum_{k=1}^3 \boldsymbol{x}_0[k] P_{q_k}^{-1}$. Given $\boldsymbol{x}_0$ and $\boldsymbol{x}_B$, let $\bar{\boldsymbol{x}}_B$ given by Eq. (7) be any point on the line connecting $\boldsymbol{x}_0$ through $\boldsymbol{x}_B$ to the edge of the simplex. Moving $\bar{\boldsymbol{x}}_B$ *away* from $\boldsymbol{x}_0$ along the line connecting $\boldsymbol{x}_0$ and $\boldsymbol{x}_B$ (orange segments), causes the pure states quantile functions $\bar{P}_{q_{1:3}}^{-1}$ to move from $P_{q_{1:3}}^{-1}$ *towards* P_0^{-1}. This corresponds to $\alpha \in [\alpha_0, 1]$ where α_0 is the smallest value of α such that $\bar{\boldsymbol{x}}_B$ still lies on the simplex. Conversely moving $\boldsymbol{x}_B$ *towards* $\boldsymbol{x}_0$ (blue segments) results in the pure state quantile functions moving *away* from P_0^{-1}. This corresponds to $\alpha \in [1, \alpha_m]$, where α_m is the largest value of α such that all $\bar{P}_{q_{1:3}}$ remain in the set of quantile functions $\mathcal{Q}$.

and the sequence of barycentric weights given data y_t for $t = 1, 2, \ldots, T$ where each observation is taken to be distributed as ρ_{B_t}.

Below we discuss two key characteristics that pose challenges to this modeling problem. The first is the non-uniqueness of the parameters (i.e., the pure state distributions and the barycentric weights) that yield a Wasserstein baryenctric distribution. The second relates to the complications that arise when we are provided only a single time series for learning a DWB model.

A. Uniqueness of the Parameters for Wasserstein barycenters

The issue of uniqueness refers to the fact that a Wasserstein barycenter is not described by a unique set of pure state distributions and barycentric weights. While the statement is true regardless of dimension (see Appendix A for an example), given the focus of this paper, here we examine the univariate case in some detail. Specifically, we provide a construction that illustrates an inverse-scaling relation between the family of pure state distributions and barycentric weights that yields the same Wasserstein barycentric distribution.

As shown in Fig. 3, assume we have a set of pure state distributions $\rho_{q_{1:K}}$ indexed by $k = 1, 2, \ldots, K$ with quantile functions, $P_{q_k}^{-1}$ and barycentric weights $\boldsymbol{x}_B \in \Delta^K$ which give rise to the barycentric distribution $\rho_B = B(\boldsymbol{x}_B, \rho_{q_{1:K}})$ with quantile function $P_B^{-1} = \sum_{k=1}^K \boldsymbol{x}_B[k] P_{q_k}^{-1}$ (Eq. (5)). For now, $\boldsymbol{x}_B$ is assumed to lie in the interior of the simplex and we consider below the cases where $\boldsymbol{x}_B$ is on a lower dimensional face or vertex. Let us choose another point $\boldsymbol{x}_0 \neq \boldsymbol{x}_B$ corresponding to barycentric quantile function $P_0^{-1} = \sum_{k=1}^K \boldsymbol{x}_0[k] P_{q_k}^{-1}$ (Eq. (5)). We construct a family of barycentric weights $\bar{\boldsymbol{x}}_B$ and pure state quantile functions, $\bar{P}_{q_{1:K}}^{-1}$ corresponding to distributions $\bar{\rho}_{q_{1:K}}$ such that $\rho_B = B(\bar{\boldsymbol{x}}_B, \bar{\rho}_{q_{1:K}})$, or in other words, $P_B^{-1} = \sum_{k=1}^K \bar{\boldsymbol{x}}_B[k] \bar{P}_{q_k}^{-1}$. Specifically, as seen in Fig. 3 we define $\bar{\boldsymbol{x}}_B$ to be a point on the line segment connecting $\boldsymbol{x}_0$ to the boundary of the simplex that passes through $\boldsymbol{x}_B$. That is,

$$\bar{\boldsymbol{x}}_B = \boldsymbol{x}_0 + \frac{1}{\alpha}(\boldsymbol{x}_B - \boldsymbol{x}_0). \tag{7}$$

The parameter α captures the scaling nature of this construction. When $\alpha = \infty$ we have $\bar{\boldsymbol{x}}_B = \boldsymbol{x}_0$. As α decreases, $\bar{\boldsymbol{x}}_B$ moves away from $\boldsymbol{x}_0$ along the blue segment connecting $\boldsymbol{x}_0$ and $\boldsymbol{x}_B$, ultimately crossing $\boldsymbol{x}_B$ when $\alpha = 1$. Further reducing α towards

0 moves $\bar{\boldsymbol{x}}_B$ along the orange component of the same line until some point $\alpha \in (0,1]$, where $\bar{\boldsymbol{x}}_B$ reaches the boundary. We denote this point as α_0.

With this definition of $\bar{\boldsymbol{x}}_B$ in Eq. (7) we construct quantile functions, $\bar{P}_{q_{1:K}}^{-1}$ such that $P_B^{-1} = \sum_{k=1}^{K} \bar{\boldsymbol{x}}_B[k] \bar{P}_{q_k}^{-1}$. Indeed,

$$
\begin{aligned}
P_B^{-1} &= \sum_{k=1}^{K} \boldsymbol{x}_B[k] P_{q_k}^{-1} = \sum_{k=1}^{K} \left(\alpha \bar{\boldsymbol{x}}_B[k] + (1-\alpha)\boldsymbol{x}_0[k] \right) P_{q_k}^{-1} \\
&= \alpha \sum_{k=1}^{K} \bar{\boldsymbol{x}}_B[k] P_{q_k}^{-1} + (1-\alpha) \sum_{k=1}^{K} \boldsymbol{x}_0[k] P_{q_k}^{-1} \\
&= \alpha \sum_{k=1}^{K} \bar{\boldsymbol{x}}_B[k] P_{q_k}^{-1} + (1-\alpha) P_0^{-1} \qquad \boxed{P_0^{-1} = \sum_{k=1}^{K} \boldsymbol{x}[k] P_{q_k}^{-1}} \\
&= \sum_{k=1}^{K} \left(\alpha \bar{\boldsymbol{x}}_B[k] P_{q_k}^{-1} + (1-\alpha) \bar{\boldsymbol{x}}_B[k] P_0^{-1} \right) \qquad \boxed{\sum_{k=1}^{K} \bar{\boldsymbol{x}}_B[k] = 1} \\
&= \sum_{k=1}^{K} \bar{\boldsymbol{x}}_B[k] \left(\alpha P_{q_k}^{-1} + (1-\alpha) P_0^{-1} \right) = \sum_{k=1}^{K} \bar{\boldsymbol{x}}_B[k] \bar{P}_{q_k}^{-1}
\end{aligned}
$$

with for each k,

$$
\bar{P}_{q_k}^{-1} = P_0^{-1} + \alpha(P_{q_k}^{-1} - P_0^{-1}). \tag{8}
$$

Eq. (8) bears a strong resemblance to Eq. (7), except now with reciprocal use of α. For $\alpha \in [\alpha_0, 1]$, $\bar{P}_{q_k}^{-1}$ is a convex combination of P_0^{-1} and $P_{q_k}^{-1}$ lying on the line segment connecting the two quantile functions. In this case, since the collection of monotone functions on $[0,1]$ is a convex set [21], $\bar{P}_{q_{1:K}}^{-1}$ will be valid quantile functions. However, for $\alpha > 1$, $\bar{P}_{q_k}^{-1}$ extends beyond P_B^{-1} along the line that connects P_0^{-1} to P_B^{-1}. In this case, $\bar{P}_{q_k}^{-1}$ is no longer a convex combination of P_0^{-1} and P_B^{-1} and not guaranteed to be a quantile function. We denote α_m as the maximum value of α such that $\bar{P}_{q_k}^{-1}$ is a valid quantile functions.

Thus, the sets of $\bar{\boldsymbol{x}}_B$ and $\bar{P}_{q_{1:K}}^{-1}$ corresponding to $\alpha \in [\alpha_0, \alpha_m]$ according to Eqs. (7) and (8) describe the family of parameters that yield the same Wasserstein barycentric distribution as $\boldsymbol{x}_B$ and $P_{q_{1:K}}^{-1}$. The reciprocal appearance of α in these equations captures the inverse-scaling relationship for this family of parameters. As shown in Fig. 3, the case where $\alpha \in [\alpha_0, 1]$ corresponds to the orange lines where *increasing* the distance of $\bar{\boldsymbol{x}}_B$ from $\boldsymbol{x}_0$ along the segment connecting $\boldsymbol{x}_B$ and $\boldsymbol{x}_0$, results in the pure states *decreasing* their distances to P_0^{-1} each along linear trajectories connecting the $P_{q_k}^{-1}$ to P_0^{-1}. Conversely, the case where $\alpha \in [1, \alpha_m]$ corresponds to the blue lines in Fig. 3 where *decreasing* the distance between $\bar{\boldsymbol{x}}_B$ and $\boldsymbol{x}_0$ by moving along the line that connects $\boldsymbol{x}_B$ to $\boldsymbol{x}_0$, results in pure state quantile functions $\bar{P}_{q_k}^{-1}$ that are now *increasing* their distance from P_0^{-1} by extending linearly along the ray from P_0^{-1} to $P_{q_k}^{-1}$.

Should $\boldsymbol{x}_B$ lie on a face of the K-simplex of dimension greater than one but less than K (i.e., *not* a vertex), this construction may be repeated by placing $\boldsymbol{x}_0 \neq \boldsymbol{x}_B$ in that same lower dimensional simplex. In such a case, we may also place $\boldsymbol{x}_0$ in the interior of the K dimensional simplex. More specifically, with $\boldsymbol{x}_B$ on a face and $\boldsymbol{x}_0$ located in the interior, it is clear that $\alpha_0 = 1$ implying that the construction above holds only if $\alpha_m > 1$. This is also the case in the event that $\boldsymbol{x}_B$ is a vertex for any value of $\boldsymbol{x}_0 \neq \boldsymbol{x}_B$. With the constraint that $\bar{P}_{q_{1:K}}^{-1}$ must remain valid quantile functions, it is possible to construct an example such that $\alpha_m = 1$, which when combined with the aforementioned case where $\alpha_0 = 1$ means that we cannot employ this construction to show non-uniqueness. However, as discussed in Appendix A neither can we conclude that the barycenter is in fact unique.

This inverse-scaling relation implies that for models such as the DWB which require learning the both pure state quantile functions and the barycentric weights corresponding to one or a sequence of Wasserstein barycentric distribution, introducing constraints on one set of parameters will have an impact on the other. We take this effect into account in Sec. IV-B when constructing regularizers to impose desirable properties on the latent state sequence and pure state distributions.

B. Model Sampling

A second characteristic of learning a DWB model relates to how we incorporate observed data into the model. Ideally, we would directly observe the data distribution ρ_{B_t} as specified by Eq. (6) at each point in time or generate an estimate of these distributions from multiple realizations of a time-series. Unfortunately, in all the practical cases of interest to us, only a single instance of the time-series is available for processing. Thus, we consider the problem of estimating the time-varying data distribution ρ_{B_t} from a single time-series that is sampled from ρ_{B_t}. To do this, we consider a window of n samples centered at t, compiled into a vector $\mathbf{y}_t = [y_{(t-\frac{n}{2})}, \ldots, y_{(t+\frac{n}{2}-1)}]^{\mathsf{T}}$. For convenience of notation, we assume n to be even. We estimate the data distribution with the distribution ρ_{y_t} based on this sample window,

$$
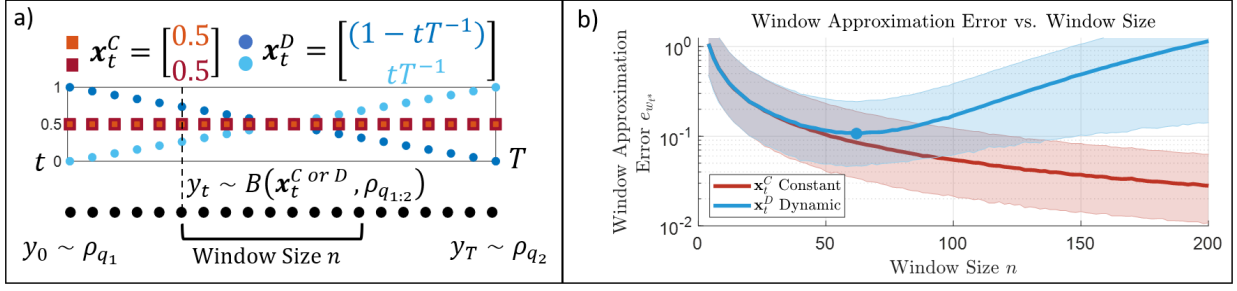\rho_{y_t} = \frac{1}{n} \sum_{i=1}^{n} \delta_{\mathbf{y}_t[i]}. \tag{9}
$$

Fig. 4: **Approximating a distribution with a window of samples**. (a) Two simulated configurations, one where $\boldsymbol{x}_t^C = [0.5, 0.5]^\mathsf{T}$ is constant, and one where $\boldsymbol{x}_t^D = [(1 - tT^{-1}), tT^{-1}]^\mathsf{T}$ is dynamically evolving. We generate a time-series $y_t \sim B(\boldsymbol{x}_t^{C \text{ or } D}, \rho_{q_{1:2}})$ where $\rho_{q_1} = \mathcal{N}(0, 5)$, and $\rho_{q_2} = \mathcal{N}(10, 0.2)^1$. (b) We estimate the distribution data distribution $\rho_{B_{t^*}}$ where $t^* = 0.5(T + 1)$, with a distribution $\rho_{y_{t^*}}$ comprised of a window of n samples centered at t^* according to Eq. (9). Varying the window size n we plot the average value and the 25/75-th quantile bands of $e_{w_{t^*}} = \mathcal{W}_2^2(\rho_{y_{t^*}}, \rho_{B_{t^*}})$ from 10^4 simulations. In the constant state case where $\rho_{y_t^*}$ consists of n IID samples from $\rho_{B_t^*}$ the average error monotonically decreases. However, in the dynamic case, where the samples in the window are independent but not identically distributed, the U-shape curve highlights the window size tradeoff where n_0 indicates the optimal window size.
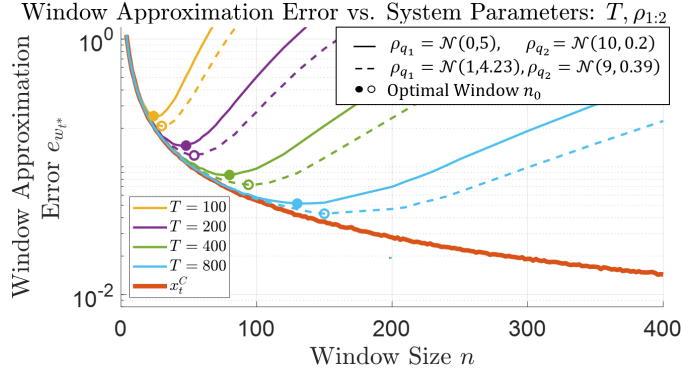

Fig. 5: **Impact of system dynamics and pure state distributions on WAE**. The average WAE is plotted for various configurations where solid lines refer to $\rho_{q_1} = \mathcal{N}(0, 5)$, $\rho_{q_2} = \mathcal{N}(10, 0.2)$ and dashed lines corresponds to a system where the Wasserstein distance between $\rho_{q_{1:2}}$ is decreased, $\rho_{q_1} = \mathcal{N}(1, 4.23)$ $\rho_{q_2} = \mathcal{N}(9, 0.39)$. In both cases, $\rho_{B_{t^*}} = \mathcal{N}(5, 1.8)$. Decreasing the per-sample change in distribution by increasing T or decreasing $\mathcal{W}_2^2(\rho_{q_1}, \rho_{q_2})$ (solid $\rightarrow$ dashed) results in smaller $e_{w_{t^*}}$. The impact of these changes on $e_{w_{t^*}}$ is greater for larger windows, which increases the optimal window n_0 corresponding to the minimum of the U-curves.

Here, $\delta_{\mathbf{y}_t[i]}$ is the Dirac-delta measure located at $\mathbf{y}_t[i]$. We then define the *window approximation error* (WAE) as,

$$e_{w_t} = \mathcal{W}_2^2(\rho_{y_t}, \rho_{B_t}). \tag{10}$$

Since the samples that constitute ρ_{y_t} are random, the WAE is a random quantity.

Eq. (9) is an empirical measure when the samples in the window are drawn IID which is only possible when the barycentric weights $\boldsymbol{x}_t$, and thus ρ_{B_t}, is constant in time. In this case, as $n \to \infty$, it is known that the distribution function of ρ_{y_t} converges almost surely to the distribution function of ρ_{B_t}, and the expected value of e_{w_t} converges to zero [22]. The complication in our case comes from the fact that $\boldsymbol{x}_t$ is changing with time, meaning that the samples in the window are no longer identically distributed as the data distribution changes from sample to sample.

This impact of a dynamically evolving distribution on the WAE is dependent on many factors including the size of the window n as well as properties of the system that impact the manner in which the data distribution changes between samples. To simplify matters, we consider a simple yet informative example in Fig. 4a of a system with two pure states and two latent state configurations, one where the latent state is *constant* with $\boldsymbol{x}_t^C = [0.5, 0.5]^\mathsf{T}$, and one where the latent state is *dynamic*, changing at a constant rate where $\boldsymbol{x}_t^D = [(1 - tT^{-1}), tT^{-1}]^\mathsf{T}$, for $t = 0, 1, \ldots, T$. For two distributions $\rho_{q_{1:2}}$, we generate a time-series by sampling independently $y_t^{C \text{ or } D} \sim B(\boldsymbol{x}_t^{C \text{ or } D}, \rho_{q_{1:2}})$. Thus, in the constant state case, y_t^C are IID. Let $\rho_{y_{t^*}}$ be the empirically estimated distribution constructed according to Eq. (9) at $t^* = 0.5(T + 1)$, the half-way point of the transition. In both cases of $\boldsymbol{x}_t^C$ and $\boldsymbol{x}_t^D$ the distribution at this time is $\rho_{B_{t^*}} = B([0.5, 0.5]^\mathsf{T}, \rho_{q_{1:2}})$.

When the window size is small, the distributional change over the sample window is also small. Therefore as seen in the left side of Fig. 4b, the dynamic case closely approximates the constant state where the average WAE decreases with respect to the

---

$^1\mathcal{N}(\mu, \sigma^2)$: Gaussian distribution with mean μ, and variance σ^2.

window size [22]. When the window size is large, the samples being included in the window are further from t^* and thus the distributions of these samples increasingly diverge from $\rho_{B_{t^*}}$. Shown by the rising and expanded quantile bands on the right side of the plot, this effect causes the average and variability of $e_{w_{t^*}}$ to increase for large windows. The resulting U-shape curve of the WAE highlights the window size tradeoff where the optimal window size n_0 corresponding to the minimum of this curve balances the benefits of having more samples for robust estimation of a distribution with the effects of using samples farther from the point of interest.

One major factor that impacts this tradeoff is the magnitude of the distributional change from sample to sample. There are number of factors that can decrease (resp. increase) the magnitude of this per-sample change in distribution including (1) decreasing (increasing) the rate of change of the system's continuous time dynamics; (2) increasing (decreasing) the sampling rate of the sensor which provides discrete measurements; or (3) decreasing (increasing) the Wasserstein distance between the pure state distributions which affects the total amount of distributional change during the transition among pure states. Continuing with the example, we demonstrate the impact of these factors on the window size tradeoff and optimal window size. Again, modeling the system as moving from ρ_{q_1} to ρ_{q_2} with dynamics specified by $\boldsymbol{x}_t^D$, factors (1) and (2) dictate the number of samples over which this transition occurs[2], thus their combined impact can be understood by varying T. To understand the impact of (3), we simulate two different pure state configurations varying $\mathcal{W}_2^2(\rho_{q_1}, \rho_{q_2})$.

The results in Fig. 5 confirm that for a given window size, decreasing the per-sample change in distribution by increasing T, or by decreasing $\mathcal{W}_2^2(\rho_{q_1}, \rho_{q_2})$ results in a decrease in the average WAE. Additionally, the increasing difference between the U-curves in Fig. 5 as we move towards larger windows confirm that the decreasing the per-sample change in distribution has an increasing benefit for larger windows where the dynamics of the system have a larger impact on the accuracy of the window estimate. This shifts the balance of the window size tradeoff as seen by the minimums of these U-curves moving to the right, implying that decreasing the per-sample change in distribution using any of the three methods mentioned increases the optimal window size. Indeed, in the limiting case where either $T \to \infty$ or $\mathcal{W}_2^2(\rho_{q_1}, \rho_{q_2}) \to 0$ in which case $\rho_{q_1} = \rho_{q_2} = \rho_{B_{t^*}}$, the samples will be drawn IID from a constant distribution. In this limit, the dynamic case converges to the constant case where the optimal window size $n_0 \to \infty$ and $e_{w_t} \to 0$.

IV. Non-Parametric and Regularized Dynamical Wasserstein Barycenters

In this section, we detail our proposed variational approach for learning the pure state distributions and barycentric weights which are the latent states of the univariate DWB model. We discuss our regularization framework that ensures that the latent state evolves smoothly over time while taking into account its effect on the pure states through the inverse-scaling relationship. We also discuss our non-parametric representation for the pure state distributions using a discrete approximation to the pure state quantile function. Finally, we detail how this leads to a least-squares formulation of the variational DWB objective and propose an algorithm for learning the parameters of the model.

A. Variational Approach to DWB Model Estimation

Training a DWB model entails estimating the pure states distributions and latent states sequence to minimize a cost function that encourages both fidelity to the data as well as model parameters that conform to prior knowledge we may have concerning the general behavior of time-series. Building on the problem formulation in Sec. III-B, instead of estimating the distribution at each point in time t, we create a sequence of N sample windows of length n that estimate the distribution of the time-series at select points. Given t_i for $i = 1, \dots, N$ as the starting index for these sample windows, let $\mathbf{y}_i = [y_{t_i}, \dots, y_{t_i+n}]^\mathsf{T}$ be the vector of samples and ρ_{y_i} the distribution according to Eq. (9) corresponding to this window of samples. Using the WAE in Eq. (10) summed over i as a data fidelity term and encoding prior information in regularizers the details of which are discussed below, the variational problem we seek to solve is,

$$\hat{\rho}_{q_{1:K}}, \hat{\boldsymbol{x}}_{1:N} = \operatorname*{argmin}_{\rho_{q_{1:K}}, \boldsymbol{x}_{1:N}}$$

$$\underbrace{\sum_{i=1}^N \mathcal{W}_2^2(\rho_{y_i}, \rho_{B_i})}_{\text{Data Fit}} + \lambda_x \underbrace{R_x(\boldsymbol{x}_{1:N})}_{\text{x-regularization}} + \lambda_q N \underbrace{R_q(\rho_{q_{1:K}})}_{\text{q-regularization}}, \tag{11}$$

where $\rho_{B_i} = B(\boldsymbol{x}_i, \rho_{q_{1:K}})$. Here, $\lambda_x, \lambda_q \geq 0$ are regularization weights. We multiply R_q by N so that it scales with length of the time-series along with the other terms in the cost function.

As highlighted by the inverse-scaling relation in Sec. III-A, any regularizer of either the latent state or pure state must consider its effect on the other parameter. Motivated by applications where the system evolves gradually over time, we propose a regularization scheme that imposes a gradually evolving latent state while ensuring that the learned pure state distributions accurately reflect the distribution of the data corresponding to when the system resides in a pure state.

B. Regularization Framework for DWB

As seen in our windowing simulations in Sec. III-B, even in a simple case, estimating a dynamically evolving distribution with a window of samples is a challenging problem. The ambiguities identified in our experiments in Sec. III-B may cause the learned latent state to vary greatly even if the system is constant or gradually evolving. Therefore, to limit its variability, we

---

[2] The rate of change (1) may have units change in distribution per second, and the sampling rate (2) has units samples per second. Thus $\frac{(1)}{(2)}$ will have units change in distribution per sample.

propose a regularizer that penalizes the squared distance between successive latent states $d^2(\boldsymbol{x}_i, \boldsymbol{x}_{i+1})$. Since $\boldsymbol{x}_i \in \Delta^K$, we consider distance measures well suited to the simplex [16]. While many such simplex distance may be chosen, for the purpose of this work, we choose this distance to be the Bhattacharrya-arccos [17]. Thus, this regularizer penalizes the total length of the latent state trajectory on the simplex according this distance metric,

$$
\begin{aligned}
R_x(\boldsymbol{x}_{1:N}) &= \sum_{i=1}^{N-1} d^2(\boldsymbol{x}_i, \boldsymbol{x}_{i+1}) \\
&= \sum_{i=1}^{N-1} \left(\arccos \left(\sum_{k=1}^{K} \sqrt{\boldsymbol{x}_i[k]\boldsymbol{x}_{i+1}[k]} \right) \right).
\end{aligned} \tag{12}
$$

With the above choice for regularizing the latent state sequence, our choice for regularization of the pure state distributions is motivated by Eq. (7) and the inverse-scaling nature of barycentric non-uniqueness. Considering our non-uniqueness construction in Eq. (7), if we set the reference point to be the current value of the latent state $\boldsymbol{x}_0 = \boldsymbol{x}_t$, and constructed point as the next state $\bar{\boldsymbol{x}}_B = \boldsymbol{x}_{t+1}$, Eq. (7) takes the form $\boldsymbol{x}_{t+1} = \boldsymbol{x}_t + \frac{1}{\alpha} \mathbf{d}_t$ where $\mathbf{d}_t$ is a vector in the simplex along which we move the latent state. Since $\mathbf{d}_t$ is essentially arbitrary, taking $\alpha > 1$ will encourage small changes in $\boldsymbol{x}_t$ as required by Eq. (12). Referring to Fig. 3, we see that in this $\alpha > 1$ regime, barycentric non-uniqueness manifests in the divergence of the quantile functions; i.e., motion along the blue line segments. That this in fact can occur is verified in our experiments in Sec. V-A1. Now, for time series that reside in each pure state at some period in time, the observed data during those periods will be representative of each of the pure state distribution. Therefore, having estimated pure state distributions diverge is undesirable if we want to accurately learn these quantities. To counteract this diverging behavior, we propose a regularizer in the space of quantile functions that penalizes the Wasserstein distance of the pure state quantile functions $P_{q_k}^{-1}$ from a reference quantile function P_0^{-1}. Here, we choose P_0^{-1} be the quantile function of $\rho_0 = B(\boldsymbol{x}_0, \rho_{q_{1:K}})$ where $\boldsymbol{x}_0 = \frac{1}{K}\mathbb{1}_K$ (with $\mathbb{1}_K$ the length K vector of ones) lies at the centroid of the simplex:

$$
\begin{aligned}
R_q(\rho_{q_{1:K}}) &= \sum_{k=1:K} \mathcal{W}_2^2 \left(\rho_{q_k}, B\left(\frac{1}{K}\mathbb{1}_K, \rho_{q_{1:K}} \right) \right) \\
&= \sum_{k=1:K} \int_0^1 \left(P_{q_k}^{-1}(\xi) - \frac{1}{K} \sum_{j=1:K} P_{q_j}^{-1}(\xi) \right)^2 d\xi,
\end{aligned} \tag{13}
$$

The two regularizers just discussed are designed to work in tandem to mitigate the effect of the mismatch between the model's Wasserstein barycenter quantile function and the quantile function corresponding to the sample window, a mismatch discussed in Sec. III-B rooted in the finite length of the sample window and the dynamics of the system. Here, R_x ensures that the the latent state evolves gradually over the simplex while R_q ensures that the pure state quantile functions do not diverge in the ways predicted by the inverse-scaling analysis in Sec. III-A. Through our simulations in Sec. V-A, we demonstrate how by appropriately balancing the regularization weights we can stabilize the estimates of the parameters of systems that can be described as gradually evolving among the system's pure states.

C. Discrete Quantile Approximation

Given the fact that the Wasserstein distance and barycenter for univariate distributions outlined in Sec. II are defined according to the quantile functions, we choose to discuss this univariate DWB model in terms of the quantile functions of the pure state distributions and their Wasserstein barycenter. Practically estimating the pure state quantile functions $P_{q_{1:K}}^{-1}$ requires a finite dimensional approximation to these functions. Toward that end we introduce the following two quantities:

Definition 1. *[n-DQA] Given a quantile function $P^{-1} : [0,1] \to \mathbb{R}$, the n-point discrete quantile approximation (n-DQA) is a monotone step function $P_n^{-1}(\xi) = P^{-1}\left(\frac{\lceil \xi n \rceil - 0.5}{n} \right)$ that is obtained by sampling the $\{\frac{0.5}{n}, \frac{1.5}{n}, ..., \frac{n-0.5}{n}\}$-th quantiles of P^{-1}.*

Definition 2. *[n-DQV] The n-point discrete quantile vector (n-DQV) $\mathbf{q} \in \mathbb{R}^n$ is comprised of the sampled quantiles from an n-point DQA: $\mathbf{q}[j] = P^{-1}\left(\frac{j-0.5}{n} \right), j = 1, 2, \cdots, n.$*

The first two plots in the top row of Fig. 6 illustrate these definitions. For a univariate distribution, the n-DQA approximates the quantile function with a monotone step function whose constant values are sampled from the quantile function on a uniform interval. Since the quantile function is monotone, these sampled quantiles and consequently the n-DQV is sorted in ascending order.

With this, we approximate the quantile functions of the pure state distributions $P_{q_{1:K}}^{-1}$ with their respective n-DQA $P_{n,q_{1:K}}^{-1}$, parameterizing them according to their n-DQV, which are denoted as $\mathbf{q}_{1:K}$. Learning these n-DQVs amounts to estimating the $\{\frac{0.5}{n}, \frac{1.5}{n}, \ldots, \frac{n-0.5}{n}\}$-quantiles of each pure state distribution. We also use this discrete quantile approach to estimate the quantile functions of the model Wasserstein barycenter $\rho_{B_i} = B(\boldsymbol{x}_i, \rho_{q_{1:K}})$. From Eq. (5) we see that the ξ-th quantile of the Wasserstein barycenter is a weighted combination of the ξ-th quantiles of $\rho_{q_{1:K}}$, with barycentric weight $\boldsymbol{x}_i$. Since the $\mathbf{q}_{1:K}$

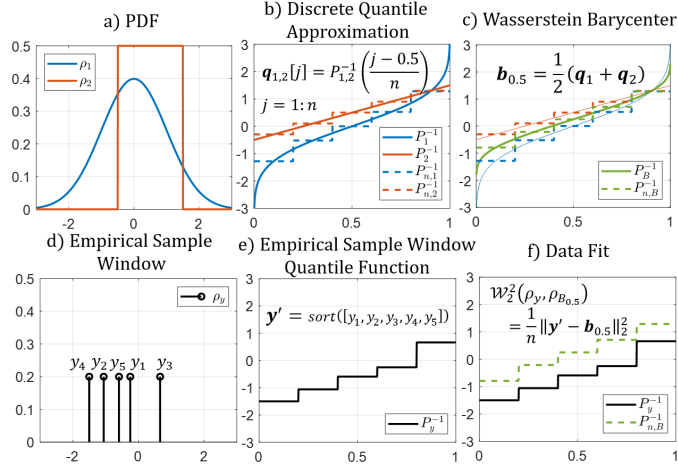Fig. 6: **Discrete quantile representation for pure states, Wasserstein barycenter, and empirical distribution function.** (a), PDF of ρ_1, ρ_2. (b) Respective quantile function and n-DQA where $n = 5$ with corresponding n-DQV $\mathbf{q}_{1:2}$. (c) Quantile function of the Wasserstein barycenter $\rho_{B_{t*}} = B([0.5, 0.5]^\top, \rho_{1:2})$ and its n-DQV, $\boldsymbol{b}_{t*}$, a weighted average of $\mathbf{q}_1$ and $\mathbf{q}_2$. (d) Distribution according to Eq. (9) corresponding to sample window with $n = 5$ points. (e), corresponding empirical quantile function P_y^{-1} where $\mathbf{y}'$ is the sorted vector of samples. (f) Since $\rho_{B_{t*}}$ and P_y^{-1} are monotone step functions sharing the same set of discontinuities, from Eq. (3) $\mathcal{W}_2^2(\rho_y, \rho_{n,B}) = \frac{1}{n}\|\mathbf{y}' - \boldsymbol{b}_{0.5}\|_2^2$

samples the quantile function of each pure state distribution at the same quantile values, the n-DQV of ρ_{B_i} is

$$\boldsymbol{b}_i = \sum_{k=1}^{K} \boldsymbol{x}_i[k]\mathbf{q}_k. \tag{14}$$

We denote corresponding n-DQA as P_{n,B_i}^{-1} and the corresponding distribution as ρ_{n,B_i}.

Using this n-DQV representation, by intentionally choosing n, the discretization level for the n-DQV, to be equal to the size of the sample window, we are able to pose a least-squares cost that approximates the data fit term in Eq. (11). To see this, we start by noting that all n-DQAs including $P_{n,q_k}^{-1}, P_{n,B_i}^{-1}$ are monotone step functions on $[0, 1]$ with discontinuities at $\{\frac{1}{n}, \frac{2}{n}, ..., \frac{n-1}{n}\}$. Additionally, shown in the bottom row of Fig. 6, the distributions ρ_{y_i} corresponding to sample windows $\mathbf{y}_i$ are discrete distributions comprised of Dirac-delta measures supported on a set of n points with uniform weights (Eq. (9)). By setting this window size n to be the same as the level of discretization used for the n-DQAs that approximate the pure state quantile functions, the quantile functions $P_{y_i}^{-1}$ corresponding to ρ_{y_i} will have the same monotone piecewise constant structure with the same set of discontinuities as the n-DQA (bottom right of Fig. 6). With $\mathbf{y}_i'$ being the vector obtained by sorting the elements of $\mathbf{y}_i$ in increasing order, the Wasserstein distance between the ρ_{y_i} and ρ_{n,B_i} is simply

$$\mathcal{W}_2^2(\rho_{y_i}, \rho_{n,B_i}) = \frac{1}{n}\|\mathbf{y}_i' - \boldsymbol{b}_i\|_2^2, \tag{15}$$

where n is size of the sample window corresponding to ρ_{y_i} and hence the length of $\mathbf{y}_i'$, as well as the level of discretization used for the n-DQA of the pure state distributions, and hence the length of $\boldsymbol{b}_i$.

D. Model Estimation and Algorithm

Using this discrete quantile parameterization for the pure state distributions allows us to pose an approximation to the variational objective function in Eq. (11) for the DWB model as a constrained nonlinear least squares problem. Let $\mathbf{Q} = [\mathbf{q}_1, ..., \mathbf{q}_K] \in \mathbb{R}^{n \times K}$ denote the matrix whose columns correspond to the n-DQV of each of the pure states, $\mathbf{Y} = [\mathbf{y}_1', ..., \mathbf{y}_N'] \in \mathbb{R}^{n \times N}$ the matrix whose columns correspond the sorted sample windows from the observed time series and $\boldsymbol{X} = [\boldsymbol{x}_1, ..., \boldsymbol{x}_N] \in \mathbb{R}^{K \times N}$ the matrix whose columns correspond to the latent state vectors across time. Pulling together the regularizers discussed in Sec. IV-B, we pose the following constrained optimization problem,

$$\hat{\mathbf{Q}}, \hat{\boldsymbol{X}} = \underset{\mathbf{Q}, \boldsymbol{X}}{\operatorname{argmin}} F(\mathbf{Q}, \boldsymbol{X})$$

$$= \underset{\mathbf{Q}, \boldsymbol{X}}{\operatorname{argmin}} \underbrace{\frac{1}{n}\|\mathbf{Y} - \mathbf{Q}\boldsymbol{X}\|_F^2}_{\text{Data fit}} + \lambda_x \underbrace{\sum_{i=1}^{N-1} d^2(\boldsymbol{X}[:, i+1], \boldsymbol{X}[:, i])}_{\text{Eq. (12)}}$$

$$+ \lambda_q \underbrace{\frac{N}{n}\left\|\mathbf{Q}\left(\mathbf{I} - \frac{1}{K}\mathbb{1}_K\mathbb{1}_K^\top\right)\right\|_F^2}_{\text{Eq. (13)}} \tag{16}$$
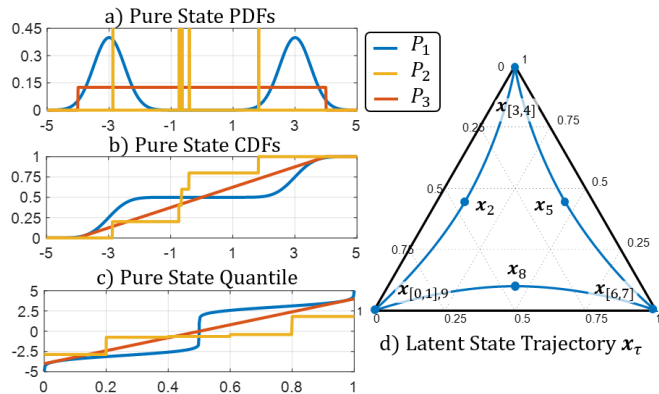
Fig. 7: **Simulated pure state distributions and latent state process.** (a) PDF, (b) CDF, and (c) quantile function for the pure state distributions used in the simulated experiments. (d) Simulated latent state trajectory as the system transitions among its three pure states.

subject to:

$$\mathbf{Q}[j+1,k] - \mathbf{Q}[j,k] > 0, \quad \begin{array}{l} k = 1, \ldots, K, \\ j = 1, \ldots, (n-1) \end{array} \tag{17}$$

$$\sum_{k=1}^{K} \boldsymbol{X}[k,i] = 1 \qquad i = 1:N \tag{18}$$

$$\boldsymbol{X}[k,i] \geq 0 \qquad \begin{array}{l} k = 1, \ldots, K \\ i = 1, \ldots, N \end{array} \tag{19}$$

The learned parameters of $\hat{\mathbf{Q}}$ and $\hat{\boldsymbol{X}}$ both have convex constraints given by Eqs. (17)-(19). The n-DQV, and thus the columns of $\mathbf{Q}$ must be sorted in ascending order (Eq. (17)). Similarly, the barycentric weights, and thus the columns of $\boldsymbol{X}$, are constrained to the set of positive matrices with rows that sum to one (Eqs. (18), (19)).

As detailed in Alg. 1, we minimize Eq. (16) using a coordinate descent approach, alternating between optimizing for $\mathbf{Q}$ and $\boldsymbol{X}$ while holding the other fixed. For both constrained optimization problems, we utilize the *sequential least squares programming* (SLSQP) optimizer [23] with python's `scipy` library [24].

Regarding initialization of the parameters $\hat{\mathbf{Q}}$ and $\hat{\boldsymbol{X}}$, we initialize the pure state distributions based on clustering the observation sample windows. Starting in line 14 in Alg. 1, we compute the similarity graph using the exponential of the negative square Wasserstein distance between the distributions estimated from any two sample windows and use spectral clustering from python's `sklearn` [25] to learn K clusters. Denoting $\mathcal{C}_k$ as the index set of sample windows that belong to cluster k, in line 22 we initialize the n-DQV of each pure state to be, $\mathbf{q}_k = \frac{1}{|\mathcal{C}_k|} \sum_{i \in \mathcal{C}_k} \mathbf{y}'_i$. We initialize the latent state $\hat{\boldsymbol{X}}[t,k] = \frac{1}{K} \forall i = 1, \ldots, N, \; k = 1, \ldots, K$, to be at the centroid of the simplex for all points in time.[3]

V. MODEL EVALUATION

Analyzing simulated and real world human activity data where the system gradually evolves among its pure states, we empirically demonstrate (1) how the proposed regularizers allows the DWB model to accurately recover the system parameters taking into account the inverse-scaling relationship between the model parameters, (2) the impact of the window size on the accuracy of estimating the system parameters of a dynamically evolving time-series, and (3) how our non-parametric formulation of the DWB problem is able to accurately learn the pure state distributions.

A. Simulated Experiments

Our approach proposes regularizers to the DWB model with the assumption that the latent state of the system transitions gradually among its pure states. Here, we consider a simulated system whose parameters reflect these properties, denoting the ground truth pure state distributions and latent states as $\rho_{q_{1:K}}$, $\boldsymbol{x}_{1:T}$ and the model estimated parameters as $\hat{\rho}_{q_{1:K}}$, $\hat{\boldsymbol{x}}_{1:T}$.

Consider a system that consists of $K = 3$ pure states with ground truth distributions,

$$\rho_{q_1} = 0.5\,\mathcal{N}(3, 0.25) + 0.25\,\mathcal{N}(-3, 0.25) \tag{20}$$

$$\rho_{q_2} = \mathcal{U}[-4, 4]$$

$$\rho_{q_3} = \frac{1}{5} \sum_{i=1}^{5} \mathcal{N}\left(\mathbf{q}[i], 1e^{-8}\right)$$

$$\mathbf{q} = [-2.88, -0.74, -0.64, -0.41, 1.82].$$

[3]Code repository: https://github.com/kevin-c-cheng/DWB_Nonparametric

---

Algorithm 1: Non-parametric and Regularized DWB

---

1 Input:
2 $y_1, ..., y_t, ..., y_T$: Time-series
3 $t_1, ..., t_i, ..., t_N$: Starting indices for sample windows
4 K: Number of pure states

5 Hyperparameters:
6 n: Sample window size η: Convergence threshold
7 λ_x, λ_q: Regularization weights

8 Output:
9 $\hat{\mathbf{Q}} \in \mathbb{R}^{n \times K}$: Stacked pure state DQVs
10 $\hat{\boldsymbol{X}} \in \mathbb{R}^{K \times N}$ Stacked barycentric weights

11 for $i = 1 : N$ **do**
12 $\quad \mathbf{y}'_i = sort(y_{t_i}, ..., y_{t_i+n})$
13 end
14 for $i = 1 : N$ **do**   // Window affinity matrix
15 $\quad$ **for** $j = 1 : N$ **do**
16 $\quad\quad \mathbf{A}[i,j] = \frac{1}{n}\left\|\mathbf{y}'_i - \mathbf{y}'_j\right\|_2^2$
17 $\quad$ **end**
18 end
19 $c_{1:N} = SpectralClustering(K, \exp(-\mathbf{A}))$;   // Cluster sample windows
20 for $k = 1, ..., K$ **do**
21 $\quad \mathcal{C}_k = \{i : c_i = k\}$
22 $\quad \mathbf{q}_k = \frac{1}{|\mathcal{C}_k|}\sum_{i \in \mathcal{C}_k} \mathbf{y}'_i$;   // Init. n-DQV
23 end
24 $\mathbf{Y} = [\mathbf{y}'_1, ..., \mathbf{y}'_N]$;   // Stacked windows
25 $\hat{\boldsymbol{X}}^{(0)} = \frac{1}{K}\mathbb{1}_K\mathbb{1}_N^{\mathsf{T}}$;   // Initialize $\boldsymbol{X}$
26 $\hat{\mathbf{Q}}^{(0)} = [\mathbf{q}_1, ..., \mathbf{q}_K]$;   // Stacked n-DQVs
27 do
28 $\quad \hat{\boldsymbol{X}}^{(i+1)} = \operatorname{argmin}_{\boldsymbol{X}} F\left(\hat{\mathbf{Q}}^{(i)}, \hat{\boldsymbol{X}}^{(i)}\right)$;   // SLSQP
29 $\quad \hat{\mathbf{Q}}^{(i+1)} = \operatorname{argmin}_{\mathbf{Q}} F\left(\hat{\mathbf{Q}}^{(i)}, \hat{\boldsymbol{X}}^{(i+1)}\right)$;   // SLSQP
30 while $F\left(\hat{\mathbf{Q}}^{(i)}, \hat{\boldsymbol{X}}^{(i)}\right) - F\left(\hat{\mathbf{Q}}^{(i+1)}, \hat{\boldsymbol{X}}^{(i+1)}\right) > \eta$;

Here, $\mathcal{U}[a,b]$ denotes a uniform distribution on the interval $[a,b]$. We show the PDF in Fig. 7a, CDF in Fig. 7b, and quantile functions in Fig. 7c for each of these distributions

Now let us consider a time-series that transitions among these three pure states following the trajectory outlined in Fig. 7d. The continuous-time latent state $\boldsymbol{x}_\tau$ alternates between pausing at each pure state for 1 second, and then transitioning to another pure state for 2 seconds moving from $\rho_1 \ldots \rho_2 \ldots \rho_3 \ldots \rho_1$ over the course of continuous time $\tau = [0, 9]$ seconds. To emulate the simulated setup in Sec. III-B of varying the per-sample change in distribution, we vary r (Hz), the rate at which we sample this continuous-time sequence to generate the ground truth latent state $\boldsymbol{x}_t$ for $t = 1, \ldots, T$. Time series are independently sampled $y_t \sim \rho_{B_t} = B(\boldsymbol{x}_t, \rho_{q_{1:3}})$ by uniformly sampling a quantile $\xi \in [0, 1]$ and evaluating the quantile function $P_{B_t}^{-1}(\xi) = \sum_{k=1}^{3} \boldsymbol{x}_t[k] P_{q_k}^{-1}(\xi)$.

From this time series, we generate a sequence of sample windows using sliding window[26], spacing out the windows on a constant interval, $t_{i+1} = t_i + \delta$. For these experiments, we choose to set $\delta = n$, which partitions the time series into a sequence of $N = \lfloor \frac{T}{n} \rfloor$ disjoint windows of size n.

With knowledge of ground truth, we can assess our model using the average distance of our learned parameters to these ground truth values according to,

$$e_q = \frac{1}{K}\sum_{k=1}^{K} \mathcal{W}_2^2\left(\rho_{q_k}, \hat{\rho}_{q_k}\right) \quad e_x = \frac{1}{N}\sum_{t=1}^{N} \|\boldsymbol{x}_t - \hat{\boldsymbol{x}}_t\|_2^2. \tag{21}$$

Since our model is unsupervised, there is no guarantee that the indexing of the learned pure states will match that of the ground truth. Therefore, when assessing our model, we assume that learned pure states (and subsequently the latent state vector) are reordered in a manner that minimizes Eq. (21).
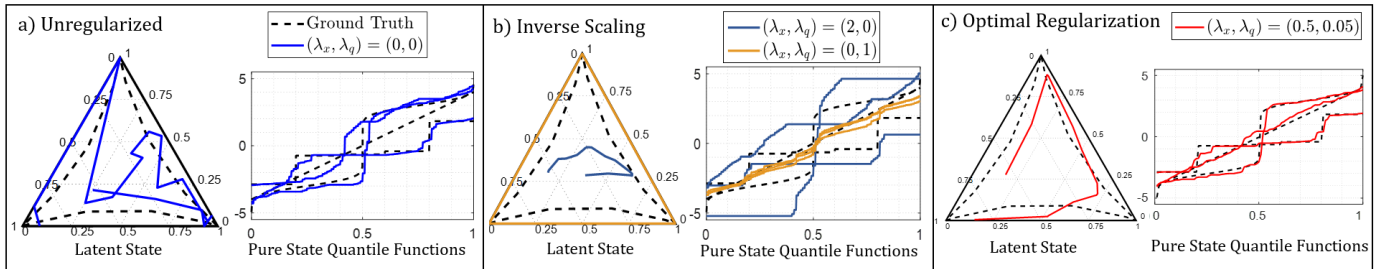
Fig. 8: **Impact of regularization and the inverse-scaling relationship**. (a) Without regularization, though the ground truth latent state evolves gradually, the learned latent state varies over the simplex. (b) Introducing just the latent state regularizer (blue) imposes some level of smoothness to the latent state, but consequently results in the pure state quantile functions diverging from their ground truth values. Conversely, including just the pure state regularizer (orange) causes the pure state quantile functions to move closer together and the latent states to move towards the boundary of the simplex. (c) Using regularization weights that minimize $e_x + e_q$ balance the effects of these two regularizers to accurately recover the gradually evolving latent state sequence and the pure states quantile functions of the time series. Plots shown for $r = 200, n = 100$.

1) Regularization and Inverse-Scaling

We demonstrate the inverse scaling relationship between the model parameters through the interaction between the latent state and pure state regularizers. As discussed in Sec. III-B, the difficulty in accurately estimating the model parameters stems from the problem of estimating dynamically evolving data distribution with a window of finite length. In Fig. 8, we see that in the absence of any regularization, the learned latent state can vary over the simplex, even when the ground truth latent state is stationary or gradually evolving. The blue lines of Fig. 8b show that introducing only the latent state regularizer that penalizes $d^2(\boldsymbol{x}_t, \boldsymbol{x}_{t+1})$ imposes some level of smoothness on the latent state. However, as discussed in Sec. IV-B this regularizer has an additional effect that is similar to the $\alpha > 1$ regime corresponding to the blue lines in Fig. 3 where the latent states trajectory contracts towards a point on the simplex. In this case, as specified by the inverse-scaling relationship, seen from the blue quantile plots in Fig. 8b, the pure state quantile functions diverge from the ground truth values. On the other hand, as seen in orange lines of Fig. 8b, having only the pure state regularizer causes the pure state quantile functions to be pulled closer together where by the inverse-scaling relationship, the latent states move away from the centroid of the simplex towards the boundary.

Only through the combination of these regularization terms can we recover the ground truth parameters of this simulated system that gradually evolves among its pure states. Using our knowledge of ground truth we perform a grid search by varying the regularization weights $\lambda_x, \lambda_q = \{1e^{-4}, 2e^{-4}, 5e^{-4}, 1e^{-3}, ..., 1e^{-1}\}$ and picking the pair that minimizes the total ground truth error $e_x + e_q$. As seen from Fig. 8c, the DWB model corresponding to these optimal regularization weights more accurately recover the ground truth parameters of this simulated system. In Sec. V-B, we parameter selection in the absences of ground truth using L-surfaces [27].

2) Impact of Window Size

In Sec. III-B, we highlighted the effect of window size and the rate of change in the data distribution on the accuracy of estimating a dynamically evolving distribution using a window of samples. Here, we seek to understand how these factors impact the accuracy of the learned DWB parameters. Using the aforementioned simulated setup, we model systems with a different rate of changes in distribution by varying the sampling rate $r = \{150, 200\}$ Hz, generating 500 time series per value of r. We then run our DWB model varying the window size between $n = 50, 75, ..., 450$. We find the optimal regularization weights performing a grid search over the range of values as Sec. V-A1. For each value of n and r, we set λ_x, λ_q to be the average of the results from this grid search performed for 5 randomly selected time series.

In this experiment we remove potential confounding variables across the various window sizes. First we use the same initialization for the pure states distributions, specifically the one generated from the spectral clustering method discussed in Sec. IV-D for the configuration of $n = 200$. Secondly, we ensure that the sequence of sample windows estimate the time-series at the same points in time, such that regardless of r, n the windows are centered at $\tau = [1.5, 2.0, ..., 7.5]$.

The results shown in Fig. 9 imply that the factors discussed in Sec. III-B that impact the ability of a sample window to accurately estimate a dynamically evolving data distribution similarly affect the ability of the DWB model to accurately learn a system's pure states and latent states from these estimated windows. The U-shape curves in Fig. 9 of the average ground truth error as a function of window size bear strong similarities to the U-shape curves illustrating the window size tradeoff in Fig. 5 where small windows lack the samples to precisely estimate the data distribution and the accuracy of large windows suffer due to the dynamics of the systems. Furthermore, we see the same trend as in Sec. III-B where decreasing the per-sample change in distribution, shown here by increasing the sampling rate r, improves the accuracy for models with larger windows shifting this window size tradeoff and the "optimal" window size towards larger windows.
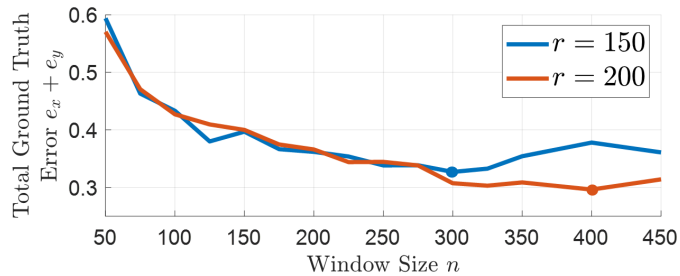
Fig. 9: **Impact of window size on simulated model accuracy**. Ground truth error averaged over 500 generated time series as a function of window size n for sampling rates of $r = 150, 200$ Hz. The U-shape plots imply that the factors that impact the ability of small and large windows to estimate a dynamically evolving data distribution similarly impact the ground truth error of the learned DWB model parameters. Additionally, increasing the sampling rate r of the time series, which results in a smaller per-sample change in the data distribution, similarly improves the model accuracy for larger windows and shifts the minimum of the U-curve towards larger windows.

B. Real World Data

In this section, we compare the performance of the non-parametric DWB model to the Gaussian DWB model [1] on univariate data. To highlight the difference between the Gaussian and non-parametric discrete quantile parameterizations, we use regularization framework proposed in this paper for both models. We evaluate on our simulated data and two human activity datasets.

1) *Beep Test (BT, proprietary):* Subjects run between two points to a metronome with increasing frequency, alternating between two states: running and standing. We use the z-axis of the 3-axis accelerometer sampled at 100 Hz.

2) *Microsoft Research Human Activity (MSR, [28]):* 126 subjects perform exercises in a gym setting. Exercises vary among subjects covering strength, cardio, cross-fit, and static exercises. Each time-series is truncated to five minutes. Discrete labels corresponding to activities are provided, thus we set K to the number of labeled discrete states in the truncated time-series (range: $K = 2$ to 7). We use the x-axis of the 3-axis accelerometer sampled at 50 Hz.

3) *Simulated Data (Sim):* Following the data generating process outlined in Sec. V-A, we simulate 500 time-series setting $r = 200$.

Evaluation: Since ground truth is not known in the real world setting, we assess performance by considering the model fit to the data, computed using the data-fit term in Eq. (11),

$$e_y = \frac{1}{N} \sum_{t=1}^{N} \mathcal{W}_2^2 \left(\rho_{y_t}, \rho_{B_t} \right) \tag{22}$$

In the non-parametric model, this distance is computed according to Eq. (15). In the Gaussian DWB model where ρ_{B_t} is Gaussian, this distance is computed using a Monte-Carlo method using $1e^5$ IID samples from ρ_{B_t}.
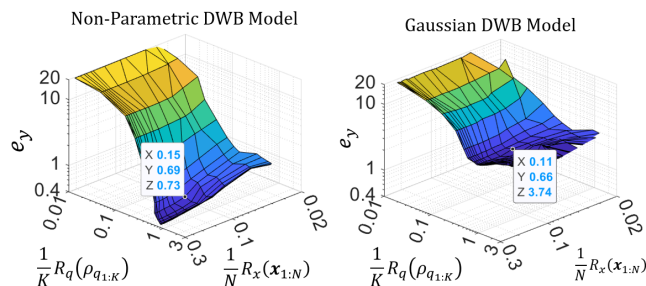

Fig. 10: **L-surface for selection of regularization weights**. L-surface for MSR time-series for (left) non-parametric and (right) Gaussian DWB models varying λ_x, λ_q from $[1e^{-5}, 1]$. Each figure plots magnitude of the data loss against the magnitude of the two regularization parameters on a log scale. Selecting the regularization weight according to L-surfaces amounts finding the λ_x, λ_q corresponding to the "corner" of the surface plot, the point where further increasing λ_x, λ_q (which corresponds to moving toward the back corner of the plot, decreaseing $\frac{1}{K} R_q(\rho_{q_{1:K}})$ and $\frac{1}{N} R_x(\boldsymbol{x}_{1:T})$) results in a sharp increase in the magnitude of the data loss.

To select regularization weights when ground truth is not available, we draw from the field of inverse problems and use L-surfaces [27]. Fig. 10 shows the L-surface for one MSR dataset plotting e_y against the magnitude of the R_x and R_q on a log scale while varying λ_x, λ_q ranging from $1e^{-5}$ to 1. For each model configuration and dataset, we use this L-surface method to pick a set of regularization parameters based on one representative time-series and apply those parameters to the rest of the dataset. These parameters are detailed in Tab. I.
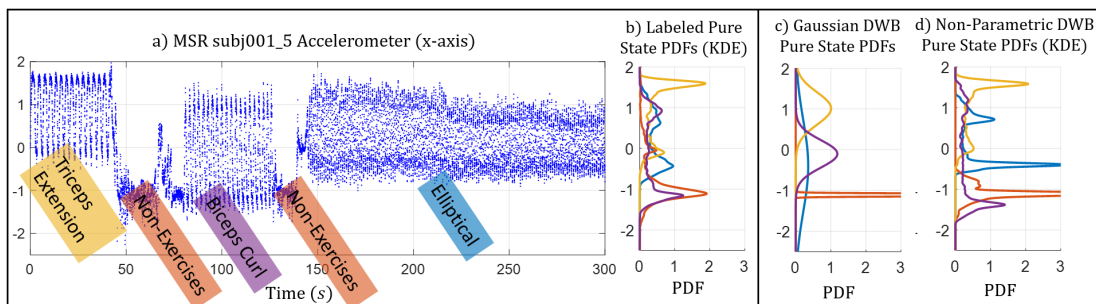
Fig. 11: **Non-parametric vs Gaussian pure states distributions for MSR data**. (a) MSR time-series consisting of 4 activities over 5 minutes. (b) Pure state distributions estimated from discrete labeled data converted to a PDF using KDE with a Gaussian kernel ($\sigma = 0.04$). (c) Learned pure states for Gaussian model from [1]. (d) Learned pure states for our proposed non-parametric model. The n-DQV $\mathbf{q}_k$ for each pure state is converted to a PDF using KDE with the same Gaussian kernel centered on the values of $\mathbf{q}_k$. The distribution of the data in each pure state activity (blue: Elliptical, yellow: Tricep Extension, purple:Bicep Curl, Red: Non-exercise) better captured using the non-parametric approach $e_y = 0.73$ compared to the Gaussian $e_y = 3.74$.

	MSR	BT	Sim
n	250	100	100
NP λ_x, λ_q	$5e^{-2}, 5e^{-3}$	$2e^{-1}, 5e^{-2}$	$2e^{-1}1e^{-2}$
Gauss λ_x, λ_q	$5e^{-2}, 2e^{-3}$	$2e^{-1}, 2e^{-2}$	$1e^{-1}2e^{-3}$

TABLE I: **DWB Model configuration**. Window size and regularization weights for non-parametric (NP) and Gaussian (Gauss) DWB models for simulated and real world datasets.

Fig. 11 compares the learned pure states of the Gaussian DWB and non-parmetric DWB for one example MSR time-series. As seen from Fig. 11b, the distribution of the data corresponding to many of the activities are clearly multi-modal and therefore not Gaussian. Compared to the Gaussian DWB approach shown in Fig. 11c, the pure states learned using the proposed non-parametric representation shown in Fig. 11d more closely match the estimated pure states from the data. This also reflected quantitatively in Tab. II where the values of e_y show that the our non-parametric DWB model better approximates the sample windows of the time series compared to the Gaussian DWB model for each of the evaluated datasets.

We note that the magnitude of the e_y, and thus the results in Tab. II are dependent on the choice of the regularization parameters, which are chosen in a partially subjective manner. However, we can see in the L-surfaces in Fig. 10 that compared to the Gaussian model the error is significantly lower in the non-parametric model. Therefore any small subjective changes in the choice of regularization parameter would not significantly alter these conclusions.

VI. Conclusions and Future Work

In this work, we build upon the DWB model for time-series in [1] by highlighting the uniqueness in learning a DWB model through demonstrating the inverse-scaling relationship between the model's pure state and latent state parameters and by discussing the challenges of estimating the data distribution of a dynamically evolving time series. We then propose a regularization framework that imposes temporal smoothness on the latent state while taking into account the impact of constraining one model parameter on the other through the inverse-scaling relationship. Finally, we show how to move beyond the over-restrictive Gaussianity assumption of [1] by using a discrete approximation to the pure state quantile function results in a least-square problem that approximates the DWB objective function in Eq. (11). Using simulated data, we demonstrate how the two proposed regularization terms work together to achieve the desired smoothness in the time evolution of the latent state as well as the impact of window size on the accuracy of the learned model. In the real world setting of human activity analysis, we demonstrate how compared to the original Gaussian DWB model our non-parametric DWB model better characterizes the time-varying data distribution of the time-series.

One limitation of the work presented in this paper is that it only applies to the univariate case where the Wasserstein distance and barycenter. To extend this model to higher dimensions, one approach is to model the pure state distributions using point clouds, the sum of uniformly weighted Dirac-delta functions supported on a set n points in $\mathbb{R}^d$. One consideration for such a model is that the Wasserstein barycenter of higher dimensional point clouds may not be unique [14]. This case is further discussed, and algorithms are proposed for computing the Wasserstein barycenter in the works of [29], [30].

	MSR	BT	Sim
NP	**1.32**	**2.07**	**2.19**
Gauss	3.76	6.00	12.33

TABLE II: **Quantitative comparison of DWB model**. Average e_y across all time-series for MSR, BT, and Sim datasets show clear benefits of the non-parametric DWB model.

Finally, in this work, we also explore the issue of approximating the data distribution of a dynamically evolving system from a window of samples by demonstrating the tradeoff between the errors associated with large and small windows. We use simulated data to succinctly illustrate the effect of window size and the system parameters that dictate the per-sample change in the distribution on the accuracy of this window estimate. Future work may consider an analytical approach to explore these factors, perhaps through the analysis of collections of independent but non-identically distributed random variables [31].

REFERENCES

[1] K. Cheng, S. Aeron, M. C. Hughes, and E. L. Miller, "Dynamical Wasserstein Barycenters for Time-series Modeling," *Advances in Neural Information Processing Systems*, vol. 34, pp. 27 991–28 003, 2021.

[2] D. Mukhin, A. Gavrilov, A. Feigin, E. Loskutov, and J. Kurths, "Principal nonlinear dynamical modes of climate variability," *Scientific reports*, vol. 5, no. 1, pp. 1–11, 2015, publisher: Nature Publishing Group.

[3] S. A. Imtiaz, "A systematic review of sensing technologies for wearable sleep staging," *Sensors*, vol. 21, no. 5, p. 1562, 2021, publisher: MDPI.

[4] Y. Chi, T. Yang, and P. Zhang, "Dynamical Mode Recognition of Triple Flickering Buoyant Diffusion Flames: from Physical Space to Phase Space and to Wasserstein Space," *arXiv preprint arXiv:2201.01085*, 2022.

[5] M. C. Hughes and E. B. Sudderth, "Nonparametric discovery of activity patterns from video collections," in *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. IEEE, 2012, pp. 25–32.

[6] O. D. Lara and M. A. Labrador, "A survey on human activity recognition using wearable sensors," *IEEE communications surveys & tutorials*, vol. 15, no. 3, pp. 1192–1209, 2012, publisher: IEEE.

[7] R. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Trans. ASME, J. Basic Eng.*, vol. 82, pp. 34–45, 1960.

[8] M. I. Ribeiro, "Kalman and extended kalman filters: Concept, derivation and properties," *Institute for Systems and Robotics*, vol. 43, p. 46, 2004.

[9] E. Fox, E. Sudderth, M. Jordan, and A. Willsky, "Nonparametric Bayesian learning of switching linear dynamical systems," *Advances in neural information processing systems*, vol. 21, 2008.

[10] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, pp. 257–286, 1989, publisher: Ieee.

[11] Z. Ghahramani and M. Jordan, "Factorial hidden Markov models," *Advances in Neural Information Processing Systems*, vol. 8, 1995.

[12] E. B. Fox, M. C. Hughes, E. B. Sudderth, and M. I. Jordan, "Joint modeling of multiple time series via the beta process with application to motion capture segmentation," *The Annals of Applied Statistics*, vol. 8, no. 3, pp. 1281–1313, 2014, publisher: Institute of Mathematical Statistics.

[13] R. J. McCann, "A convexity principle for interacting gases," *Advances in mathematics*, vol. 128, no. 1, pp. 153–179, 1997, publisher: Elsevier.

[14] M. Agueh and G. Carlier, "Barycenters in the Wasserstein space," *SIAM Journal on Mathematical Analysis*, vol. 43, no. 2, pp. 904–924, 2011, publisher: SIAM.

[15] T.-M. Nguyen and S. Volkov, "On a class of random walks in simplexes," *Journal of Applied Probability*, vol. 57, no. 2, pp. 409–428, 2020, publisher: Cambridge University Press.

[16] J. Martín-Fernández, C. Barceló-Vidal, V. Pawlowsky-Glahn, A. Buccianti, G. Nardi, and R. Potenza, "Measures of difference for compositional data and hierarchical clustering methods," in *Proceedings of IAMG*, vol. 98, 1998, pp. 526–531, issue: 1.

[17] F. Åström, S. Petra, B. Schmitzer, and C. Schnörr, "Image labeling by assignment," *Journal of Mathematical Imaging and Vision*, vol. 58, no. 2, pp. 211–238, 2017, publisher: Springer.

[18] G. Peyré, M. Cuturi, and others, "Computational optimal transport: With applications to data science," *Foundations and Trends® in Machine Learning*, vol. 11, no. 5-6, pp. 355–607, 2019, publisher: Now Publishers, Inc.

[19] N. Bonneel and H. Pfister, "Sliced Wasserstein barycenter of multiple densities," 2013.

[20] P. Embrechts and M. Hofert, "A note on generalized inverses," *Mathematical Methods of Operations Research*, vol. 77, no. 3, pp. 423–432, 2013, publisher: Springer.

[21] R. Rakestraw, "The convex cone of n-monotone functions," *Pacific Journal of Mathematics*, vol. 43, no. 3, pp. 735–752, 1972, publisher: Mathematical Sciences Publishers.

[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, publisher: Springer.

[23] D. Kraft, "A software package for sequential quadratic programming," *Forschungsbericht- Deutsche Forschungs- und Versuchsanstalt fur Luft- und Raumfahrt*, 1988.

[24] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, and others, "SciPy 1.0: fundamental algorithms for scientific computing in Python," *Nature methods*, vol. 17, no. 3, pp. 261–272, 2020, publisher: Nature Publishing Group.

[25] L. Buitinck, G. Louppe, M. Blondel, F. Pedregosa, A. Mueller, O. Grisel, V. Niculae, P. Prettenhofer, A. Gramfort, J. Grobler, and others, "API design for machine learning software: experiences from the scikit-learn project," in *European Conference on Machine Learning and Principles and Practices of Knowledge Discovery in Databases*, 2013.

[26] S. Aminikhanghahi and D. J. Cook, "A survey of methods for time series change point detection," *Knowledge and information systems*, vol. 51, no. 2, pp. 339–367, 2017, publisher: Springer.

[27] D. H. Brooks, G. F. Ahmad, R. S. MacLeod, and G. M. Maratos, "Inverse electrocardiography by simultaneous imposition of multiple constraints," *IEEE Transactions on Biomedical Engineering*, vol. 46, no. 1, pp. 3–18, 1999, publisher: IEEE.

[28] D. Morris, T. S. Saponas, A. Guillory, and I. Kelner, "RecoFit: using a wearable sensor to find, recognize, and count repetitive exercises," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2014, pp. 3225–3234.

[29] M. Cuturi and A. Doucet, "Fast computation of Wasserstein barycenters," in *International conference on machine learning*. PMLR, 2014, pp. 685–693.

[30] E. Anderes, S. Borgwardt, and J. Miller, "Discrete Wasserstein barycenters: Optimal transport for discrete data," *Mathematical Methods of Operations Research*, vol. 84, no. 2, pp. 389–409, 2016, publisher: Springer.

[31] H. A. David and H. N. Nagaraja, *Order statistics*. John Wiley & Sons, 2004.
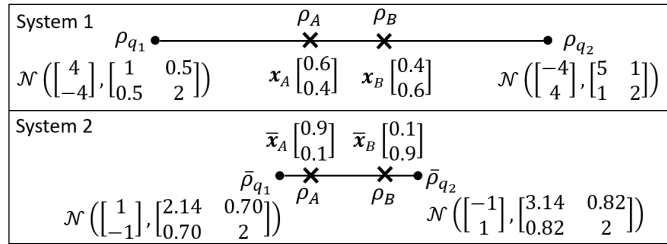
Fig. 12: **Non-uniqueness of Wasserstein barycenter parameters for multivariate Gaussians**. Both systems of pure state distributions and barycentric weights result in the same barycentric distributions, $\rho_A = B(\boldsymbol{x}_A, \rho_{q_{1:2}}) = B(\bar{\boldsymbol{x}}_A, \bar{\rho}_{q_{1:2}})$ and $\rho_B = B(\boldsymbol{x}_B, \rho_{q_{1:2}}) = B(\bar{\boldsymbol{x}}_B, \bar{\rho}_{q_{1:2}})$. Consistent with the inverse-scaling relationship discussed in Sec. III-A, compared to system 1, in system 2 the Wasserstein distance between $\bar{\rho}_{q_1}$ and $\bar{\rho}_{q_2}$ is diminished while the distance between $\bar{\boldsymbol{x}}_A$ and $\bar{\boldsymbol{x}}_B$ is increased.

APPENDIX A
EXTENSIONS OF DWB NON-UNIQUENESS

While the discussion of this paper pertains mainly to the univariate case, the issue of uniqueness discussed in Sec. III also exists in the multivariate case. Consider the two dimensional systems specified in Fig. 12 of systems with two Gaussian pure states. For both systems, $\rho_A = B(\boldsymbol{x}_A, \rho_{q_{1:2}}) = B(\bar{\boldsymbol{x}}_A, \bar{\rho}_{q_{1:2}})$ and $\rho_B = B(\boldsymbol{x}_B, \rho_{q_{1:2}}) = B(\bar{\boldsymbol{x}}_B, \bar{\rho}_{q_{1:2}})$. This multivariate example also exhibits the inverse scaling relationship discussed in Sec. III-A. The Wasserstein distance between the pure states in system 1 (ρ_{q_1}, ρ_{q_2}) is *larger* than that of system 2 ($\bar{\rho}_{q_1}, \bar{\rho}_{q_2}$), however, the resulting distance between the latent states corresponding to ρ_A, ρ_B in system 1 ($\boldsymbol{x}_A, \boldsymbol{x}_B$) is *smaller* than that of system 2 ($\bar{\boldsymbol{x}}_A, \bar{\boldsymbol{x}}_B$).

Furthermore, we can create an example where the construction specified in Sec. III-A does not result in multiple possible values of $\boldsymbol{x}_B$ and $P_{q_{1:K}}^{-1}$ and thus cannot be used to demonstrate the non-uniquness of the Wasserstein barycenter parameters.

Let $\rho_{q_1} = \delta_0$ and $\rho_{q_2} = U[0,1]$, where $P_{q_1}^{-1}(\xi) = 0$ and $P_{q_2}^{-1}(\xi) = \xi$ for $\xi \in [0,1]$. Then the Wasserstein barycenter $\rho_t = B(\boldsymbol{x}_t, \rho_{1:2})$ for any barycentric weight $\boldsymbol{x}_t = [(1-t), t]^\mathsf{T}$ will have distribution $U[0, t]$ and quantile function $P_t^{-1}(\xi) = t\xi$. Per the construction in Sec. III-A, let us pick $\boldsymbol{x}_B = [1, 0]$ to be at a vertex, thus $P_B^{-1}(\xi) = 0$ and let $\boldsymbol{x}_0 = [0.5, 0.5]^\mathsf{T}$, thus in $P_0^{-1}(\xi) = 0.5\xi$. According to Eq. (7), $\bar{\boldsymbol{x}}_B$ falls of the simplex for any $\alpha < 1$ thus making $\alpha_0 = 1$. Similarly, according to Eq. (8), $\bar{P}_B^{-1}(\xi) = (1-\alpha)t\xi$, which breaks the monotonically increasing constraint of quantile functions for any $\alpha > 1$, thus restricting $\alpha_m = 1$. Thus, according to the construction in Sec. III-A, the set $[\alpha_0, \alpha_m] = 1$.

Although this example breaks the argument of the non-uniqueness of the Wasserstein barycenter parameters for this specific construction used to highlight the inverse-scaling relationship, it does not mean that the parameters of the Wasserstein barycenter are indeed unique. For example, since $\boldsymbol{x}_B$ is taken to be at a vertex $[1, 0]^\mathsf{T}$ resulting in $\rho_B = \rho_{q_1}$, any valid distribution can be chosen for ρ_{q_2} and still have $\rho_B = B([1, 0]^\mathsf{T}, \rho_{q_{1:2}})$.