

THE WASSERSTEIN DISTANCES BETWEEN PUSHED-FORWARD MEASURES WITH APPLICATIONS TO UNCERTAINTY QUANTIFICATION*

AMIR SAGIV[†]

Abstract. In the study of dynamical and physical systems, the input parameters are often uncertain or randomly distributed according to a measure ϱ . The system’s response f pushes forward ϱ to a new measure $f_*\varrho$ which we would like to study. However, we might not have access to f , but to its approximation g . This problem is common in the use of surrogate models for numerical uncertainty quantification (UQ). We thus arrive at a fundamental question – if f and g are close in an L^q space, does the measure $g_*\varrho$ approximate $f_*\varrho$ well, and in what sense? Previously, it was demonstrated that the answer to this question might be negative when posed in terms of the L^p distance between probability density functions (PDF). Instead, we show in this paper that the Wasserstein metric is the proper framework for this question. For domains in \mathbb{R}^d , we bound the Wasserstein distance $W_p(f_*\varrho, g_*\varrho)$ from above by $\|f - g\|_q$. Furthermore, we prove lower bounds for the cases where $p = 1$ and $p = 2$ (for $d = 1$) in terms of moments approximation. From a numerical analysis standpoint, since the Wasserstein distance is related to the cumulative distribution function (CDF), we show that the latter is well approximated by methods such as spline interpolation and generalized polynomial chaos (gPC).

Keywords. Wasserstein; Uncertainty-Quantification; Density-Estimation; Optimal Transport; Approximation.

AMS subject classifications. 28A10; 60A10; 65D99.

1. Introduction

1.1. Problem formulation. Suppose a domain $\Omega \subseteq \mathbb{R}^d$ is equipped with a Borel probability measure ϱ and that a function $f: \Omega \rightarrow \mathbb{R}$ pushes forward ϱ to a new measure $\mu := f_*\varrho$, i.e., $f_*\mu(B) = \varrho(f^{-1}(B))$ for every Borel set $B \subseteq \mathbb{R}$. We wish to characterize μ , but only have access to a function g which approximates f . If $\|f - g\|_{L^q(\Omega, \varrho)}$ is small, does $\nu := g_*\varrho$ approximate μ well, and if so in what sense?

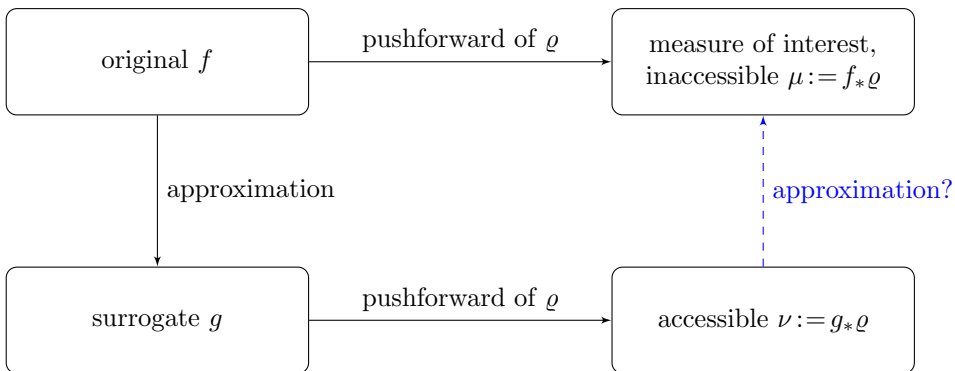


FIG. 1.1. *The schematic structure of the problem. If $\|f - g\|_p$ is small, how close are μ and ν ? In other words, is the dashed arrow “justified”?*

*Received: March 23, 2019; Accepted (in revised form): November 18, 2019. Communicated by Siddhartha Mishra.

[†]Department of Applied Mathematics, Tel Aviv University, Tel Aviv 6997801, Israel (asagiv88@gmail.com; as6011@columbia.edu).

1.2. Motivation. To motivate this rather abstract question, consider the following toy example: a harmonic oscillator is described by the ordinary differential equation (ODE) $y''(t) + y = 0$ with $y(0) = 0$ and $y'(0) = v$. Suppose we are interested in $f(v) = y^2(\pi/2; v)$. By solving this ODE, we know that $f(v) = [v \sin(\pi/2)]^2 = v^2$. In many other cases, however, we do not have direct access to f , but only to its approximation g . This could happen for various reasons – it may be that we can only compute $f(\alpha)$ numerically, or that we approximate f using an asymptotic method. Following on the harmonic oscillator example, suppose we know $f(v)$ only at four given points v_1, v_2, v_3 , and v_4 . For any other value of v , we approximate $f(v)$ by $g(v)$, which linearly interpolates the adjacent values of f , see Figure 1.2(a).

The parameters and inputs of physical systems are often noisy or uncertain. We thus assume in the harmonic oscillator example that the initial speed v is drawn uniformly at random from $[1, 2]$. In these settings, $f(v)$ is random, and we are interested in the *distribution* of $f(v)$ over many experiments. Even though f and g look similar in Figure 1.2(a), the probability density functions (PDF) of $\mu = f_*\varrho$ and $\nu = g_*\varrho$, denoted by p_μ and p_ν respectively, are quite different, see Figure 1.2(b). We would therefore like to have guarantees that ν approximates the original measure of interest μ well.




FIG. 1.2. Solutions of $y''(t) + y = 0$ with $y'(0) = v$ and $y(0) = 0$. (a) $f(v) = y^2(t = \pi/2; v)$ (solid) and $g(v)$, its piecewise linear interpolant based on four exact samples (dash-dots). (b) The PDFs of $\mu = f_*\varrho$ (solid) and $\nu = g_*\varrho$ (dash-dots), where ϱ is the uniform distribution on $[1, 2]$.

It might seem obvious that the distance between f and g controls the distance between μ and ν . This hypothesis fails, however, when one estimates this distance using the PDFs p_μ and p_ν . For example, let $f(\alpha) = \alpha$ and $g(\alpha) = \alpha + \delta \sin((10\delta)^{-1}\alpha)$, where $1 \gg \delta > 0$. Since $\|f - g\|_\infty = \delta$, the two functions are seemingly indistinguishable from each other, see Figure 1.3(a). Consider the case where ϱ is the Lebesgue measure on $[0, 1]$. Then, since both functions are monotonic, $p_\mu(y) = 1/f'(f^{-1}(y)) = 1$ and $p_\nu(y) = 1/g'(g^{-1}(y))$, see [13] for details. Hence, p_ν is onto $[1.1^{-1}, 0.9^{-1}] \approx [0.91, 1.11]$ and so $\|p_\mu - p_\nu\|_\infty > 0.1$, irrespective of δ , see Figure 1.3(b). The lack of apparent correspondence between $\|f - g\|_q$ and $\|p_\mu - p_\nu\|_p$ for *any* pair of integers p and q suggests that the PDFs are not a well-suited metric for the problem depicted in Figure 1.1. *Instead, in this paper we propose the Wasserstein distance as the proper framework to measure the distance between μ and ν .*

1.3. Relevant literature. The harmonic oscillator example in Section 1.2 serves as a toy example for a broad class of problems. While the ODE $y''(t) + y = 0$ can be solved explicitly, many other differential equations do not admit such closed-form solutions. Instead, we only have an approximation for the quantities of interest at our disposal. Indeed, the general settings presented above have spurred numerous

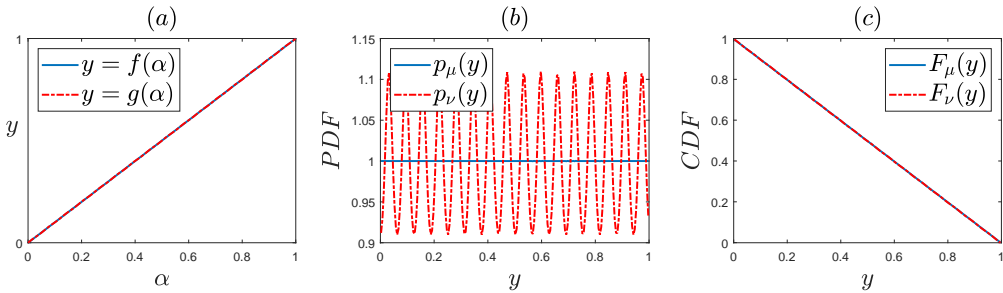


FIG. 1.3. (a) $f(\alpha) = \alpha$ (solid) and $g(\alpha) = \alpha + 10^{-3} \sin(100\alpha)$ (dash-dots). The two lines are indistinguishable. (b) The PDFs of $\mu = f_*\rho$ and $\nu = g_*\rho$, where ρ is the Lebesgue measure on $\Omega = [0, 1]$. The CDFs of the same measures. The two lines are indistinguishable.

papers in a field of computational science known as Uncertainty-Quantification (UQ), see e.g., [13, 19, 38, 48–50]. Perhaps surprisingly, the full approximation of μ (rather than its moments alone) in these particular settings received little theoretical attention in the literature, even though it is of practical importance in diverse fields such as ocean waves [1], computational fluid dynamics [8, 43], hydrology [9], aeronautics [15], biochemistry [21], and nonlinear optics [26, 32]. Even though $\|f - g\|_q$ does not control $\|p_\mu - p_\nu\|_p$ in general (see e.g., Figure 1.3), a previous result by Ditkowski, Fibich, and the author gives sufficient conditions for PDF approximation:

THEOREM 1.1 (Ditkowski, Fibich, and Sagiv [13]). *Let $f \in C^2([0, 1]^d)$ and let $g_h \in C^2([0, 1]^d)$ be an interpolant of f on a tensor grid of maximal spacing $h > 0$ such that*

$$\|f - g_h\|_\infty, \|\nabla f - \nabla g_h\|_\infty \leq Kh^\tau,$$

where $K = K(f, d)$ and $\tau > 0$ is fixed. Then

$$\|p_\mu - p_\nu\|_{L^p} \leq \tilde{K}h^\tau,$$

for every $1 \leq p < \infty$, with a constant $\tilde{K} = \tilde{K}(f, d, q)$.

The conditions on g are motivated by spline interpolation method, see Section 4 for further details. Theorem 1.1 is, to the best of our knowledge, a first result in the direction of this paper’s main question. Even so, Theorem 1.1 is limited in several ways:

- (1) The demand $|\nabla f| \geq \tau_f > 0$ is an arbitrary condition from an application standpoint.
- (2) The differentiability and the pointwise derivative-approximation conditions $\|\nabla f - \nabla g\|_\infty \lesssim h^\tau$ are strong demands which many other approximation methods do not fulfill.
- (3) It is essential that the domain Ω is compact for the proof to hold.
- (4) Even when $d = 1$, it is required that $d\rho(\alpha) = c(\alpha)d\alpha$ with $c \in C^1(\bar{\Omega})$. For comparison, absolute continuity is a weaker condition, as it requires that $c \in C(\Omega) \cap L^1(\Omega)$.

The Wasserstein distance (see Section 1.4) is thus proposed to measure the distance between μ and ν since it does not suffer from the drawbacks of the norms $\|p_\mu - p_\nu\|_p$. Admittedly, the L^p distances between the PDFs are both natural in practice and are associated with rich statistical theory; for $p = 1$, then $\|p_\mu - p_\nu\|_1$ is twice the total variation [12], and $\|p_\mu - p_\nu\|_2^2$ is the integrated square error, which is a building block in non-parametric statistics [42]. Nevertheless, the analysis of the norms $\|p_\mu - p_\nu\|_p$ in

terms of the functions f and g can be technically cumbersome; if e.g., ϱ is the Lebesgue measure, then $p_\mu(y)$ is proportional to $\int_{f^{-1}(y)} 1/|\nabla f| d\sigma$, where $d\sigma$ is the $(d-1)$ dimensional surface measure [13]. Moreover, the distance $\|p_\mu - p_\nu\|_p$ is difficult to work with since it assumes that μ and ν have distributions. This is not always the case. For example, let ϱ be the Lebesgue measure on $[0, 1]$ and let

$$f_k(\alpha) = \begin{cases} 0 & x \in [0, \frac{1}{2}], \\ (x - \frac{1}{2})^k & x \in [\frac{1}{2}, 1], \end{cases} \quad k \geq 1.$$

Although f_k is in $C^k([0, 1])$, the measure $\mu_k = (f_k)_* \varrho$ is not an absolutely-continuous measure and does not have a PDF since $\mu(\{0\}) = 1/2$. It is therefore natural to look for other ways to measure the distance between μ and ν . There are many ways to define distances between probabilities and measures, such as total variation, mutual information, and Kullback-Leibler divergence. The equivalencies and relationships between these norms, metrics, and semi-metrics are the topics of many studies, see e.g., [16].

1.4. The Wasserstein distance. In order for us to choose the proper metric between μ and ν , we revisit Figure 1.3. While the two PDFs seem very different on a local scale, they are quite similar on a coarser scale. For example, $\mu([0.3, 0.4]) \approx \nu([0.3, 0.4])$ and so, if we were to ask what is the probability that the results of many experiments are between 0.3 and 0.4, then both μ and ν would have provided similar answers. More loosely speaking, since p_ν is oscillatory, the regions where $p_\nu > p_\mu$ and the regions where $p_\nu < p_\mu$ are adjacent, and therefore cancel-out each other. The PDF, on the other hand, is the derivative of the measure, and it is therefore heavily affected by local differences. Another disadvantage of the norm $\|p_\mu - p_\nu\|_q$ is that it does not take geometry into account. Consider for example a family of standard Gaussian measures with mean $t \in \mathbb{R}$, i.e., $p_{\mu,t}(y) = \exp(-(y-t)^2)/\sqrt{2\pi}$ (see Figure 1.4). Then for every $t > 2$, $\|p_{\mu,t} - p_{\mu,0}\|_1 \approx 2$, regardless of whether $t = 3$ or $t = 10$ or $t = 1,000$.




FIG. 1.4. Gaussian distributions centered at $t=0$ (solid), $t=3$ (dash-dots), and $t=10$ (dots). Which of the latter two Gaussians is closer to the $t=0$ Gaussian in Wasserstein distance, and which in the L^q sense?

A widely-popular metric that overcomes some of the above issues is the Wasserstein metric. Given two probability measures μ and ν on \mathbb{R} with $p \geq 1$ finite moments, the Wasserstein distance of order p is defined as

$$W_p(\mu, \nu) := \left[\inf_{\gamma \in \Gamma} \int |x - y|^p d\gamma(x, y) \right]^{\frac{1}{p}}, \tag{1.1a}$$

where Γ is the set of all measures γ on \mathbb{R}^2 for which μ and ν are marginals, i.e.,

$$\mu(x) = \int_{\mathbb{R}} \gamma(x, y) dy, \quad \nu(y) = \int_{\mathbb{R}} \gamma(x, y) dx. \tag{1.1b}$$

If the p -th moments of μ and ν are finite, then a minimizer exists, $W_p(\mu, \nu)$ is finite, and it is a metric [34, 45]. Intuitively, the Wasserstein distance with $p=1$ computes the minimal work (distance times force) by which one can transfer a mound of earth that “looks” like μ to a one that “looks” like ν , and it is therefore referred to as the earth-mover’s distance.

As noted, some of the difficulties in approximating the PDFs arise from the inverse proportion between p_μ and p_ν and the *gradients* of f and g , respectively. It is therefore natural to avoid these issues by considering the integral of the PDF, the cumulative distribution function (CDF)

$$F_\sigma(y) := \sigma([y, \infty)) = \int_y^\infty p_\sigma(t) dt,$$

for any Borel measure σ . Indeed, the Wasserstein distance of order $p=1$ is related to the CDF by the following theorem.

THEOREM 1.2 (Salvemini [33], Vallender [44]). *For any two probability Borel measures μ and ν on \mathbb{R} ,*

$$W_1(\mu, \nu) = \int_{\mathbb{R}} |F_\mu(x) - F_\nu(x)| dx.$$

This theorem reinforces the notion that W_1 is not as sensitive to local effects as $\|p_\mu - p_\nu\|_p$. Indeed, Figure 1.3(c) shows that the two CDFs of μ and ν are almost indistinguishable. Furthermore, in the previous Gaussians example (see Figure 1.4), $W_1(p_{\mu,t}, p_{\mu,0}) = t$ by direct computation of the CDFs, then, and the same can be proven for $p=2$ as well [17, 24]. Hence, the geometric distance between the Gaussians matters under the Wasserstein metric. Generally, Wasserstein distances are a central object in optimal transport theory [34, 45], and have also become increasingly popular in such diverse fields as image processing [25, 31], optimization and neural networks [3], well-posedness proofs for partial differential equations with an associated gradient-flow [7], and numerical methods for conservation laws [36, 40].

1.5. Structure of the paper. The rest of the paper is organized as follows: Section 2 presents the main theoretical results of this paper. The upper bounds on W_p (Theorems 2.1 and 2.2) are presented in Section 2.1, and the lower bounds on W_1 (Corollary 2.1) and W_2 (Theorem 2.4) are presented in Section 2.2. The proofs and some technical details of these results are presented in Section 3. Finally, in Section 4 the theoretical results are applied to the numerical analysis of uncertainty quantification methods, and a numerical example is presented.

2. Main results

2.1. Upper bounds. In what follows, $\Omega \subseteq \mathbb{R}^d$ is a Borel set, ϱ is a Borel probability measure on Ω , $f, g: \Omega \rightarrow \mathbb{R}$ are measurable, $\mu = f_* \varrho$, $\nu = g_* \varrho$, and $L^p = L^p(\Omega, \varrho)$ for any $1 \leq p \leq \infty$ unless stated otherwise.

THEOREM 2.1. *Let f and g be continuous on $\bar{\Omega}$.*

(i) If $f, g \in L^\infty(\Omega, \varrho)$, then for every $p \geq 1$

$$W_p(\mu, \nu) \leq \|f - g\|_\infty.$$

(ii) If Ω is bounded and $f, g \in L^p(\Omega, \varrho)$ then

$$W_p(\mu, \nu) \leq \|f - g\|_p.$$

This result is sharp. Let ϱ be any probability measure on $[0, 1]$ and let $f(\alpha) \equiv x_0$ and $g(\alpha) \equiv y_0$, for some $x_0, y_0 \in \mathbb{R}$. Then μ and ν , are the Dirac delta distributions centered at x_0 and y_0 , respectively, and the only distribution $\gamma \in \Gamma$ is $\gamma = \delta_{(x_0, y_0)}$. Hence, $W_p^p(\mu, \nu) = |x_0 - y_0|^p = \|f - g\|_\infty^p$. Furthermore, as opposed to Theorem 1.1, this theorem does not even demand that f and g be differentiable, and puts no restrictions on the Borel measure ϱ . Though this theorem is only valid for domains in \mathbb{R}^d , a generalization of part (i) to (infinite-dimensional) Polish spaces has been achieved by Boussaid [6].

Part (ii) of Theorem 2.1 uses L^p information to bound $W_p(\mu, \nu)$. In many cases, however, upper bounds on $f - g$ are known only in a specific L^q space. The next theorem shows how L^q error estimates can provide nontrivial upper bounds on $W_p(\mu, \nu)$ for any p , even if $p \neq q$.

THEOREM 2.2. *Under the assumptions (i)+(ii) of Theorem 2.1, then for every $p, q \geq 1$,*

$$W_p(\mu, \nu) \lesssim \|f - g\|_\infty^{\frac{p}{q+p}} \|f - g\|_q^{\frac{q}{q+p}},$$

where \lesssim denotes inequality up to a constant which depends only on p and q .

This limit agrees with Theorem 2.1 in the limit $q \rightarrow \infty$ and when $q = p$ (up to a constant). Furthermore, for any $q \neq p$, the bound in Theorem 2.2 may improve the L^∞ bound in Theorem 2.1, since ϱ is a probability measure, $(f - g) \in L^\infty \cap L^q$, and so $\|f - g\|_q \leq \|f - g\|_\infty$.

2.2. Lower bounds. The W_1 lower bound is the direct result of the Monge-Kantorovich duality, see Section 3.3 for details and proof.

COROLLARY 2.1. *If $f, g \in C(\bar{\Omega})$ and Ω is bounded, then*

$$|\mathbb{E}_\varrho f - \mathbb{E}_\varrho g| \leq W_1(\mu, \nu) \leq \|f - g\|_{L^1(\Omega, \varrho)}.$$

Moreover, if $f \geq g$ almost everywhere with respect to ϱ , then

$$W_1(\mu, \nu) = \|f - g\|_{L^1(\Omega, \varrho)}.$$

We note that since the upper bound is sharp (see discussion on Theorem 2.1) and since equality might hold, the lower bound is sharp too. We further note that in the case where Ω is the unit circle, lower bounds on W_1 in terms of the Fourier coefficients of f were proved by Steinerberger [37].

Next, to bound $W_2(\mu, \nu)$ from below, we introduce two concepts: the Sobolev space \dot{H}^{-1} and the symmetric decreasing rearrangement. For any Borel measure σ on \mathbb{R} , define the semi-norm

$$\|\sigma\|_{\dot{H}^{-1}(\mathbb{R})} := \sup_{\|q\|_{\dot{H}^1(\mathbb{R})} \leq 1} |\langle q, \sigma \rangle|,$$

where $\|q\|_{\dot{H}^1}^2 = \int |q'(x)|^2 dx$ [2]. Note that $\|\sigma\|_{\dot{H}^{-1}} < \infty$ only if $\sigma(\mathbb{R}) = 0$. Another way to understand the Sobolev semi-norm \dot{H}^{-1} and to compare it to the more frequently used L^2 norm is through Fourier analysis. By Plancharel Theorem

$$\|\sigma\|_{L^2}^2 = \int_{\mathbb{R}} |\hat{\sigma}(\xi)|^2 d\xi, \quad \|\sigma\|_{\dot{H}^{-1}}^2 = \int_{\mathbb{R}} \left| \frac{\hat{\sigma}(\xi)}{|\xi|} \right|^2 d\xi,$$

where $\hat{\sigma}$ is the Fourier transform of σ [2]. Thus, if μ and ν are different only in high frequencies, then their L^2 difference might be much higher than their \dot{H}^{-1} difference (due to the $1/|\xi|$ term in the integral). Intuitively, it means that highly local effects in $\sigma = \mu - \nu$ are “subdued” in the negative Sobolev semi-norm. This is analogous to the way local effects in the PDFs are subdued in the W_1 distance, i.e., in the CDFs (see Figure 1.3). As noted, this property also characterizes the Wasserstein distance, and indeed Loeper [23] and Peyre [27] related $W_2(\mu, \nu)$ to $\|\mu - \nu\|_{\dot{H}^{-1}}$ in the following theorem:

THEOREM 2.3 (Loeper [23], Peyre [27]). *Let μ and ν be probability measures on \mathbb{R} with densities $p_\mu, p_\nu \in L^\infty(\mathbb{R})$, respectively. Then,*

$$\|\mu - \nu\|_{\dot{H}^{-1}} \leq \max\{\|p_\mu\|_\infty, \|p_\nu\|_\infty\}^{\frac{1}{2}} W_2(\mu, \nu),$$




FIG. 2.1. (a) $f(\alpha) = 5(1 + \alpha \sin(10\alpha)e^{-\alpha^2})$. (b) $f^*(\alpha)$, the symmetric decreasing rearrangement of f , with respect to the Lebesgue measure on $[0, 1]$.

We now introduce the symmetric decreasing rearrangement by an absolutely-continuous Borel probability measure on $\Omega \subseteq \mathbb{R}^d$ [22]. The symmetric decreasing rearrangement of a measurable set A is

$$A^* = \{\alpha \in \Omega \mid \varrho(B(0, 1)) \cdot |\alpha|^d \leq \varrho(A)\},$$

where $B(0, 1) \subset \mathbb{R}^d$ is the unit ball around the origin. Next, for a measurable non-negative function $f: \mathbb{R}^d \rightarrow \mathbb{R}_+$, define the symmetric decreasing rearrangement as

$$f^*(\alpha) = \int_0^\infty \mathbb{1}_{\{\alpha' \in \Omega \mid f(\alpha') > t\}^*}(r) dt, \quad r := |\alpha|,$$

where $\mathbb{1}_B$ is the identifier of a set $B \subseteq \mathbb{R}^d$. For a numerical example of the symmetric decreasing rearrangement, see Figure 2.1. In more intuitive terms, f^* is the unique monotonic decreasing function such that $\varrho(A(f, t)) = \varrho(A(f^*, t))$ for all $t \in \mathbb{R}$, where $A(f, t) := \{\alpha \text{ s.t. } f(\alpha) \geq t\}$ are the super-level sets of f . Moreover, since f^* is monotonic decreasing, one also has that $A(f^*, t)$ is the interval $[0, \varrho(A(f, t))]$. The symmetric

decreasing rearrangement is an important object in real analysis [22], with notable properties such as the Pólya-Szego inequality [28]

$$\|f\|_p = \|f^*\|_p, \quad \|\nabla f^*\|_p \leq \|\nabla f\|_p,$$

for all $p \geq 1$. Hence, the symmetric decreasing rearrangement f^* minimizes all Sobolev $W^{1,p}$ norms among the class of functions with the same super-level sets, it can be said to be the “canonical” representative of this class.

THEOREM 2.4. *Let I be a closed and bounded interval equipped with an absolutely-continuous probability measure ϱ with a bounded and continuous weight function $r(\alpha)$, i.e., $d\varrho(\alpha) = r(\alpha)d\alpha$, and let $f, g \in C^1$ with $|(f^*)'|, |(g^*)'| > \tau > 0$. Then, for every $k \in \mathbb{N}$*

$$W_2(\mu, \nu) \geq A_k |\mathbb{E}_\varrho f^k - \mathbb{E}_\varrho g^k|,$$

where A_k is a positive coefficient given by

$$A_k = A_k(f, g, r) = \frac{\sqrt{2k-1}}{k} (\max(f, g)^{2k-1} - \min(f, g)^{2k-1})^{-\frac{1}{2}} \tau^{\frac{1}{2}} \|r\|_\infty^{-\frac{1}{2}},$$

and the max and min are taken over all $x \in I$.

We remark that even though $A_k = A_k(f, g, r)$ depends on f and g , it does not depend directly on $f - g$. Hence, for a sequence $(g_n(\alpha))_{n=1}^\infty$ which converges uniformly to f , for each $k \in \mathbb{N}$, then $A_k(f, g_n, r)$ would converge to a positive constant as $n \rightarrow \infty$. A specific example of the computation of the coefficients A_k can be found in Section 3.4.

3. Proofs of main results and technical discussion

3.1. Proof of Theorem 2.1. We begin with the case where f and g are *uniformly continuous* in Ω . Let $\varepsilon > 0$, then by uniform continuity there exists $\eta = \eta(\varepsilon) > 0$ such that $|f(\alpha) - f(\beta)| < \varepsilon$ and $|g(\alpha) - g(\beta)| < \varepsilon$ for every $\alpha, \beta \in \Omega$ such that $|\alpha - \beta| < \eta$. Let $L \in \mathbb{N}$ and partition $[-L, L]^d$ to M equal-size boxes $\{\tilde{I}_j\}_{j=1}^M$ such that $\text{diam}(\tilde{I}_j) < \eta$. Let $I_j = \tilde{I}_j \cap \Omega$ for every $1 \leq j \leq M$ and let $I_{M+1} := \Omega \setminus [-L, L]^d$. Next, let

$$\mu_j := f_* \varrho|_{I_j}, \quad \nu_j := g_* \varrho|_{I_j},$$

i.e., the measures induced by $f(I_j)$ and $g(I_j)$ for every $1 \leq j \leq M+1$. Since $\int_{\mathbb{R}} \mu_j = \int_{\mathbb{R}} \nu_j = \varrho(I_j)$, we can transport μ to ν by transporting each μ_j to ν_j . Even though this might not be the optimal transport between μ and ν , since W_p is defined as an infimum over all transports, then

$$W_p^p(\mu, \nu) \leq \sum_{j=1}^M W_p^p(\mu_j, \nu_j) + W_p^p(\mu_{M+1}, \nu_{M+1}), \tag{3.1a}$$

for any $1 \leq j \leq M+1$, where

$$\begin{aligned} W_p^p(\mu_j, \nu_j) &:= \inf_{\gamma \in \Gamma_j} \int_{f(I_j) \times g(I_j)} |x - y|^p d\gamma(x, y) \\ &\leq \left(\sup_{(x, y) \in f(I_j) \times g(I_j)} |x - y|^p \right) \varrho(I_j), \end{aligned} \tag{3.1b}$$

where Γ_j is the set of all measures whose marginals are μ_j and ν_j . For $1 \leq j \leq M$, since $\text{diam}(I_j) < \eta$, then by uniform continuity for any $t \in I_j$

$$\sup_{(x,y) \in f(I_j) \times g(I_j)} |x - y|^p \leq (|f(t) - g(t)| + 2\varepsilon)^p.$$

Here the proofs of the L^p and L^∞ bounds slightly diverge and we begin with proving that $W_p(\mu, \nu) \leq \|f - g\|_\infty$. For any $1 \leq j \leq M$ then

$$(|f(t) - g(t)| + 2\varepsilon)^p \leq (\|f - g\|_\infty + 2\varepsilon)^p.$$

Similarly, for $j = M + 1$, the supremum in (3.1) is bounded from above by $(\|f\|_\infty + \|g\|_\infty)^p$. Combining these bounds together, we have that

$$W_p^p(\mu, \nu) \leq \|f - g\|_\infty^p \sum_{j=1}^M \varrho(I_j) + o(\varepsilon) \sum_{j=1}^M \varrho(I_j) + (\|f\|_\infty + \|g\|_\infty)^p \varrho(I_{M+1}).$$

Since ϱ is a probability measure $\sum_{j=1}^M \varrho(I_j) = \varrho(\Omega) = 1$ and as $L \rightarrow \infty$ the third term on the right-hand-side vanishes. Hence $W_p^p(\mu, \nu) \leq \|f - g\|_\infty^p + 2o(\varepsilon)$ for every $\varepsilon > 0$, and so $W_p \leq \|f - g\|_\infty$.

Next, consider the case where f, g are continuous on Ω , but not uniformly continuous. Then for any two sequences $\varepsilon_n \rightarrow 0$ and $L_n \rightarrow \infty$, choose $\eta_n = \eta(\varepsilon_n, L_n)$ which satisfies the uniform continuity condition on the compact domain $\Omega \cap [-L_n, L_n]$. Then, by partitioning this domain into sufficiently many boxes $M_n = M(\eta_n) = M(\varepsilon_n, L_n)$ such that $\text{diam}(I_{j,n}) \leq \eta_n$, the proof holds as $n \rightarrow \infty$.

Finally, we prove that $W_p(\mu, \nu) \leq \|f - g\|_p$. Here we require that Ω is bounded, and so we can choose L such that $\Omega \subseteq [-L, L]^d$. For $1 \leq j \leq M$ we have that, for some $t_j \in I_j$,

$$\begin{aligned} \sup_{(x,y) \in f(I_j) \times g(I_j)} |x - y|^p &\leq (|f(t_j) - g(t_j)| + 2\varepsilon)^p \\ &= |f(t_j) - g(t_j)|^p + o(\varepsilon). \end{aligned}$$

Substituting this inequality in (3.1) yields

$$W_p^p(\mu, \nu) \leq \sum_{j=1}^M |f(t_j) - g(t_j)|^p \varrho(I_j) + o(\varepsilon) \sum_{j=1}^M \varrho(I_j).$$

As the partition is refined (i.e., $M \rightarrow \infty$ and $\eta \rightarrow 0$), the first element on the right-hand-side converges to $\|f - g\|_{L^p(\varrho)}$. Since ϱ is a probability measure, $\sum_{j=1}^M \varrho(I_j) = 1$, and so the second element on the right-hand-side is $o(\varepsilon)$. Since this inequality is true for any $\varepsilon > 0$, the proof follows.

3.2. Proof of Theorem 2.2. Define $\Omega_r := \{\alpha \in \Omega \mid |f(\alpha) - g(\alpha)| \geq r\}$ for any $r > 0$, and let $\mu_{\Omega_r}, \mu_{\Omega \setminus \Omega_r}, \nu_{\Omega_r}$, and $\nu_{\Omega \setminus \Omega_r}$ be the measure induced by $f(\Omega_r), f(\Omega \setminus \Omega_r), g(\Omega_r)$, and $g(\Omega \setminus \Omega_r)$, respectively. For any $p \geq 1$,

$$W_p^p(\mu, \nu) \leq W_p^p(\mu_{\Omega_r}, \nu_{\Omega_r}) + W_p^p(\mu_{\Omega \setminus \Omega_r}, \nu_{\Omega \setminus \Omega_r}). \tag{3.2}$$

The first term on the right-hand-side of (3.2) is bounded from above by $\|f - g\|_\infty^p \varrho(\Omega_r)$, due to Theorem 2.1. To bound $\varrho(\Omega_r)$, note that

$$\|f - g\|_{L^q(\Omega)}^q \geq \|f - g\|_{L^q(\Omega_r)}^q \geq \varrho(\Omega_r) \cdot r^q,$$

where the first inequality is due to monotonicity of ϱ , and the last inequality is due to the continuity of $|f - g|^q$. Hence, $\varrho(\Omega_r) \leq \|f - g\|_q^q r^{-q}$, and so the first term in the right-hand-side of (3.2) is bounded from above by $\|f - g\|_\infty^p \|f - g\|_q^q r^{-q}$. Since the L^∞ upper bound of Theorem 2.1 is applicable to f and g , and since $\varrho(\Omega \setminus \Omega_r) \leq 1$, then the second term on the right-hand-side of (3.2) is bounded from above by $\|f - g\|_{L^\infty(\Omega \setminus \Omega_r)}^p \leq r^p$. Having bounded from above both terms on the right-hand-side of (3.2), then

$$W_p^p(\mu, \nu) \leq \|f - g\|_\infty^p \|f - g\|_q^q r^{-q} + r^p.$$

To minimize the right-hand-side of this inequality, we derive with respect to r and get that the minimum is achieved at $r_{\min} = (qp^{-1} \|f - g\|_q^q \|f - g\|_\infty^p)^{1/(p+q)}$, and so

$$\begin{aligned} W_p(\mu, \nu) &\leq [\|f - g\|_\infty^p \|f - g\|_q^q r_{\min}^{-q} + r_{\min}^p]^{\frac{1}{p}} \\ &\lesssim \|f - g\|_\infty^{\frac{p}{q+p}} \|f - g\|_q^{\frac{q}{q+p}}. \end{aligned}$$

3.3. Proof of Corollary 2.1. The Monge-Kantorovich duality states that [45]

$$W_1(\mu, \nu) = \sup \left\{ \left| \int_{\mathbb{R}} w(y) (d\mu(y) - d\nu(y)) \right| : L(w) \leq 1 \right\},$$

where $L(w)$ is the Lipschitz constant of w . So, to prove a non-trivial lower bound for $\mu = f_*\varrho$ and $\nu = g_*\varrho$, it is sufficient to provide any function w for which the integral is not zero. Let $w(y) = y$. Since $L(w) = 1$, then $W_1(\mu, \nu) \geq |\int_{\mathbb{R}} y d\mu(y) - \int_{\mathbb{R}} y d\nu(y)|$, which, by change of variables, means that $W_1(\mu, \nu) \geq |\int_{\Omega} (f(\alpha) - g(\alpha)) d\varrho(\alpha)|$. Combined with Theorem 2.1 we arrive at the corollary.

3.4. Proof of Theorem 2.4.

Proof. By definition of the symmetric decreasing rearrangement, $\mu = f_*^*\varrho$ and $\nu = g_*^*\varrho$. Moreover, since the theorem requires that $|(f^*)'|, |(g^*)'| > \tau > 0$, we can assume without loss of generality that f and g are strongly monotonically decreasing. Next, we have the following standard lemma (for proof, see e.g., [13]):

LEMMA 3.1. *Let $h \in C^1(I)$ be piecewise monotonic, let $d\varrho(\alpha) = r(\alpha)d\alpha$ where r is continuous in Ω . Then the PDF of the measure $\sigma = h_*\varrho$ is given by*

$$p_\sigma(y) = \sum_{\alpha \in h^{-1}(y)} \frac{r(h^{-1}(y))}{|h'(h^{-1}(y))|}, \quad y \in \text{range}(h).$$

Hence, by definition and the above lemma

$$\begin{aligned} \|\mu - \nu\|_{\dot{H}^{-1}} &= \sup_{\|q\|_{\dot{H}^1} \leq 1} \int_{\mathbb{R}} q(y) (p_\mu(y) - p_\nu(y)) dy \\ &= \sup_{\|q\|_{\dot{H}^1} \leq 1} \left| \int_{\mathbb{R}} q(y) \frac{r(f^{-1}(y))}{f'(f^{-1}(y))} dy - \int_{\mathbb{R}} q(y) \frac{r(g^{-1}(y))}{g'(g^{-1}(y))} dy \right|. \end{aligned}$$

Consider the first integral under the supremum. By change of variables $y = f(x)$, we have

$$\int_{\mathbb{R}} q(y) \frac{r(f^{-1}(y))}{f'(f^{-1}(y))} dy = \int_I q \circ f(x) \frac{r(x)}{f'(x)} f'(x) dx = \int_I q \circ f d\varrho(x).$$

The respective change of variable for the second integral under the supremum yields

$$\|\mu - \nu\|_{\dot{H}^{-1}} = \sup_{\|q\|_{\dot{H}^1} \leq 1} \left| \int_I (q \circ f(x) - q \circ g(x)) d\varrho(x) \right|. \tag{3.3}$$

For ease of notations, denote $M = \max_{x \in I} \{f(x), g(x)\}$ and $m = \min_{x \in I} \{f(x), g(x)\}$. Since f and g are continuous on a closed bounded interval, both m and M are finite. Fix $k \in \mathbb{N}$, and let $q_k(x) = c_k x^k$, where the normalization constant $c_k := (\sqrt{2k-1}/k)(M^{2k-1} - m^{2k-1})^{-1/2}$ is chosen so that $\|q_k\|_{\dot{H}^1[m, M]} = 1$.¹ Hence, substituting q_k in (3.3) for every $k \in \mathbb{N}$

$$\begin{aligned} \|\mu - \nu\|_{\dot{H}^{-1}} &\geq \left| \int_I (q_k \circ f(x) - q_k \circ g(x)) d\varrho(x) \right| \\ &= c_k \left| \int_I f^k(x) - g^k(x) d\varrho(x) \right| \\ &= c_k |\mathbb{E}_\varrho f^k - \mathbb{E}_\varrho g^k|. \end{aligned}$$

Finally, to bound W_2 from below we need Theorem 2.3, and so we need to compute $\|p_\mu\|_\infty$ and $\|p_\nu\|_\infty$. As noted, since $f = f^*$ is strictly decreasing, it is also continuously differentiable almost everywhere. Hence, by the result noted above, $p_\mu = r(f^{-1}(y))/|f'(f^{-1}(y))|$ almost everywhere, and so $\|p_\mu\|_\infty \leq \tau^{-1}\|r\|_\infty$. Since the same holds for g and ν as well, we substitute in the bound from Theorem 2.3 and get that

$$\begin{aligned} W_2(\mu, \nu) &\geq [\max\{\|p_\mu\|_\infty, \|p_\nu\|_\infty\}]^{-\frac{1}{2}} \|\mu - \nu\|_{\dot{H}^{-1}} \\ &\geq [\tau\|r\|_\infty^{-1}]^{\frac{1}{2}} \|\mu - \nu\|_{\dot{H}^{-1}} \\ &\geq \tau^{\frac{1}{2}} \|r\|_\infty^{-\frac{1}{2}} c_k |\mathbb{E}_\varrho f^k - \mathbb{E}_\varrho g^k|. \end{aligned}$$

□

We complement the proof by an example of a direct computation of the coefficients A_k . Let $f(\alpha) = 3\alpha - 3$, $g(\alpha) = 2\alpha - 2$ and ϱ is the Lebesgue measure on $[0, 1]$, then by direct computation we have that $M = 0$, $m = -3$, $\|r\|_\infty = 1$, $\tau = 2$, and so

$$A_k = \frac{\sqrt{2k+1}}{k} 3^{-k+\frac{1}{2}} 2^{\frac{1}{2}} \cdot 1, \quad k \in \mathbb{N}.$$

4. Convergence of uncertainty-quantification methods and numerical examples

We apply the main theoretical results of this paper to the analysis of uncertainty quantification (UQ) methods. In many applications, one can only compute the quantity of interest $f(\alpha)$ for a finite subset of α values $\{\alpha_j\}_{j=1}^N$. To compute $\mu = f_*\varrho$, we first

¹It might seem that the choice of the interval $[m, M]$ is made ad-hoc. However, this proof can be carried out in the space $\dot{H}^{-1}(\mathbb{R})$ regardless, by the following construction: extend $q_k(y)$ to \mathbb{R} by setting $q_k(y) = q_k(m)$ for $y < m$ and $q_k(y) = q_k(M)$ for $y > M$. Since outside $[m, M]$, $q'_k \equiv 0$, then $\|q_k\|_{\dot{H}^1(\mathbb{R})} = \|q_k\|_{\dot{H}^1([m, M])}$, and $\langle q_k, \mu - \nu \rangle$ is unchanged too since $\mu - \nu$ is supported only on $[m, M]$. Our choice is also consistent with Theorem 2.3, since these also “take place” on the supports of μ and σ .

use these sampled values $\{f(\alpha_j)\}_{j=1}^N$ to construct an approximate function $g(\alpha)$, and then we approximate $\mu \approx \nu = g_* \varrho$, see Figure 1.1. This measure-approximation problem is characterized by the following trade-off: The computational cost comes from direct computation of the samples $\{f(\alpha_j)\}_{j=1}^N$,² and so it *increases* linearly with N . On the other hand, we expect the approximation error to *decrease* with the sample size N , i.e., as we improve the sampling resolution. The question is, therefore, how to construct g such that μ is accurately approximated with a small sample size N .

In terms of numerical analysis, the main result of this paper is that upper bounds on $\|f - g\|_q$ *do guarantee* an upper bound on the Wasserstein distances $W_p(\mu, \nu)$. This in turn immediately implies an upper bound on the L^1 distance between the CDFs, due to the previously-noted Salvemini-Vallender identity $W_1(\mu, \nu) = \|F_\mu - F_\nu\|_1$, see Theorem 1.2.

The upper bounds on the Wasserstein-error stand in sharp contrast to the L^q errors between the PDFs, since in general an upper bound on $\|f - g\|_q$ does not guarantee an upper bound on $\|p_\mu - p_\nu\|_{L^p}$, for any finite p and q [13]. We therefore see that the way we define the approximation-error in this problem is not a mere technicality, but rather determines the results of the convergence analysis. Furthermore, we see that CDFs are “easier” to approximate than PDFs, in the sense that it is easier to guarantee their efficient approximation.

We demonstrate the applicability of our theory for two approximation methods (surrogate models), spline interpolation and generalized polynomial chaos (gPC).

4.1. Spline interpolation. Given an interval $\Omega = [\alpha_{\min}, \alpha_{\max}]$ and grid-points $\alpha_{\min} = \alpha_1 < \alpha_2 < \dots < \alpha_N = \alpha_{\max}$, an interpolating m -th order spline $g(\alpha) \in C^{m-1}(\Omega)$ is a piecewise polynomial of order m that interpolates $f(\alpha)$ at the grid-points, endowed with some additional boundary conditions so that it is unique. See [11, 29] for comprehensive expositions on splines, see [30, 35] for their extension to multidimensional domains via tensor-products, and see [4, 19] for their applicability to UQ problems. Since Theorem 4.1 is directly applicable to spline interpolation [13], if g is the spline interpolant of f , then the PDFs of μ and ν are close, i.e., $\|p_\mu - p_\nu\|_{L^p}$ is bounded from above for any $1 \leq p < \infty$. We show that in these settings, the Wasserstein distance between the measures is also bounded from above.

THEOREM 4.1. *Let $f \in C^{m+1}([0, 1]^d)$, let $g(\alpha)$ be its (tensor-product) spline interpolant of order m on a (tensor-product) grid of maximal grid size h , and let ϱ be a probability Borel measure. Then, for every $p \geq 1$,*

$$W_p(\mu, \nu) \lesssim h^{m+1} \approx N^{-\frac{m+1}{d}}, \quad \|F_\mu - F_\nu\|_1 \lesssim N^{-\frac{m+1}{d}},$$

where N is the total number of interpolation points, and where \lesssim and \approx denote inequality and equality up to constants independent of h and N , respectively.

Proof. The error of spline interpolation is controlled by the following theorem.

THEOREM 4.2 (de Boor [11] and Hall and Meyer [20]). *Let $f \in C^{m+1}([\alpha_{\min}, \alpha_{\max}])$, and let $g(\alpha)$ be its “not-a-knot”, clamped or natural m -th spline interpolant. Then*

$$\|(f(\alpha) - g(\alpha))^{(j)}\|_{L^\infty[\alpha_{\min}, \alpha_{\max}]} \leq C_{\text{spl}}^{(j)} \|f^{(m+1)}\|_\infty h^{m+1-j}, \quad j = 0, 1, \dots, m-1,$$

²Since g is given in closed form, e.g., by a polynomial, it is computationally cheap to estimate the measure $\nu = g_* \varrho$. Computing $f(\alpha_j)$, on the other hand, might involve a full numerical solution of a PDE.

where $C_{\text{spl}}^{(j)} > 0$ is a universal constant that depends only on the type of boundary condition and j, m , and $h = \max_{1 < j \leq N} |\alpha_j - \alpha_{j-1}|$.

This result is extended for higher dimensions using the the construction of tensor-product grid and tensor-product splines. The definitions here become more technical, and we refer to Schultz [35] for further details. We note that even in the multidimensional case, the error is still bounded by the spacing h^{m+1-j} . However, the number of grid points N is proportional to h^{-d} (this is the so-called curse of dimensionality which we previously mentioned). By the above error bounds, and by Theorem 2.1, we have that $W_p(\mu, \nu) \leq \|f - g\|_\infty \lesssim h^{m+1}$. \square

Theorem 4.1 is stronger than Theorem 1.1 in three aspects. First, Theorem 4.1 holds for a broader function class than the application of Theorem 1.1 to splines, since it does not require that $|\nabla f| > \tau_f > 0$, or even that the underlying measure ϱ should be absolutely continuous. Second, Theorem 4.1 is non-trivial even for those functions for which Theorem 1.1 does apply. To obtain a “trivial” upper bound, note that for any two probability measures of μ and ν with PDFs p_μ and p_ν , then

$$W_1(\mu, \nu) \leq \frac{1}{2} \text{diam}(\Omega) \cdot \|p_\mu - p_\nu\|_1,$$

where $\text{diam}(\Omega)$ is the diameter of $\Omega = \text{supp}(\mu) \cup \text{supp}(\nu)$ [16]. Since f and g are continuous on a compact set, they are bounded, and so the supports of μ and ν are bounded as well. Hence, $\text{diam}(\Omega) < \infty$, and so by Theorem 1.1, $W_1(\mu, \nu) \leq Kh^m$. Theorem 4.1, however, guarantees an additional order of accuracy and so non-trivially improves the previous results.³ Finally, Theorem 4.1 applies not only for $p=1$ but for all $p \geq 1$.




FIG. 4.1. Approximation of $\mu = f_*\varrho$ where $f(\alpha) = \alpha/2 + \tanh(9\alpha)$ and ϱ is the uniform probability measure on $[-1, 1]$. (a) $W_1(\mu, \nu)$ where g is the spline interpolant of f on a uniformly spaced grid (rectangles) and a polynomial fit $\sim N^{-4.59}$ (solid), as predicted by Theorem 4.1. (b) Same, but where g is the collocation gPC approximation of f (circles) and an exponential fit $\sim 10^{-0.7N}$ (solid), as predicted by Theorem 4.3. (c) L^1 error of the PDFs using the collocation gPC method.

Numerical example. Let

$$f(\alpha) = \frac{\alpha}{2} + \tanh(9\alpha), \quad \Omega = [-1, 1], \quad d\varrho(x) = \frac{1}{2} dx. \tag{4.1}$$

We use a cubic spline interpolant on a grid of N uniformly-spaced points, with the not-a-knot boundary condition [11]. Theorem 4.1 guarantees that in this case $W_p(\mu, \nu) \lesssim N^{-4}$. Indeed, Figure 4.1(a) shows the W_1 difference between the two measures as a function of N , and that the convergence rate is $N^{-4.59}$.

³Unfortunately, Theorem 4.1 cannot improve the L^1 bound in Theorem 1.1 since, in general, $\|p_\mu - p_\nu\|_1 \lesssim W_1(\mu, \nu)$ only for finite spaces [16].

4.2. Generalized Polynomial Chaos (gPC). Next, we turn to study W_p convergence of L^2 -spectral methods, for which PDF convergence is an open problem. We focus on the widely popular gPC.

Review of the collocation gPC method. For a more detailed exposition, see e.g., [18, 49]. Let the Jacobi polynomials $\{p_n(x)\}_{n=0}^\infty$ be the orthogonal polynomials with respect to ϱ , i.e., p_n is a polynomial of degree n , and $\int_{-1}^1 p_n(\alpha)p_m(\alpha)d\varrho(\alpha) = \delta_{n,m}$, see [39] for details. This family of orthogonal polynomials constitutes an orthonormal basis of the space $L^2(\Omega, \varrho)$, i.e., for every $f \in L^2$ one can expand

$$f(\alpha) = \sum_{n=0}^{\infty} \hat{f}(n)p_n(\alpha), \quad \hat{f}(n) := \int_{\Omega} f(\alpha)p_n(\alpha)d\varrho(\alpha).$$

This expansion converges spectrally, i.e., if f is in C^r , then $\{\hat{f}(n)\} \lesssim n^{-r}$, and if f is analytic in an ellipse $E \subseteq \mathbb{C}$ that contains $[-1, 1]$, then $|\hat{f}(n)| \lesssim e^{-\gamma n}$, for some $\gamma > 0$. Thus, one has that for such analytic functions

$$\|f - \pi_N(f)\|_2 \lesssim e^{-\gamma N}, \quad \pi_N(f) := \sum_{n=0}^N \hat{f}(n)p_n(\alpha).$$

The expansion coefficients $\{\hat{f}(n)\}$ can be approximated using the Gauss quadrature

$$\hat{f}(n) \approx \hat{f}_N(n) := \sum_{j=1}^N f(\alpha_j)p_n(\alpha_j)w_j, \quad n = 0, 1, \dots, N-1,$$

where $\{\alpha_j\}_{j=1}^N$ are the quadrature points, the distinct and real roots of $p_N(\alpha)$, and w_j are the quadrature weights [10]. We define the gPC collocation approximation g_N to be the truncated expansion of f with the quadrature-based coefficients $\hat{f}_N(n)$. We remark that this approximation method has a much simpler form – the gPC collocation approximation is also the unique interpolating polynomial of f of order $N-1$ at the quadrature points [13]. We remark that our theory can also be applied to Galerkin-gPC methods [48].

Density estimation in UQ: The main appeal of the gPC method is its spectral L^2 convergence. As noted above, it is an open question whether this can be used to prove convergence of the PDFs, i.e., an upper bound on $p_\mu - p_\nu$ in some L^p . However, Theorem 2.2 implies that spectral L^2 convergence of g_N to f can yield fast convergence of $W_p(\mu, \nu)$ for any $1 \leq p < \infty$.

THEOREM 4.3. *Let f be analytic in an ellipse in the complex plane that contains $[-1, 1]$, and let $d\varrho(\alpha) = k(1 - \alpha)^{\beta_1}(1 + \alpha)^{\beta_2}d\alpha$, for any $\beta_1, \beta_2 \in \mathbb{R}$ and a proper normalization constant $k = k(\beta_1, \beta_2)$. Let $g(\alpha)$ be the collocation gPC approximation of f , i.e., the N -th order polynomial interpolant of f at the respective Gauss quadrature points. Then, for every $p \geq 1$,*

$$W_p(\mu, \nu) \lesssim e^{-\gamma N}, \quad \|F_\mu - F_\nu\|_1 \lesssim e^{-\gamma N}, \quad n \rightarrow \infty,$$

where γ does not depend on N .

Proof. If f is analytic, the truncated expansion has the exponential accuracy

$$\|f(\alpha) - \sum_{n=0}^{N-1} \hat{f}(n)p_n(\alpha)\|_2 \lesssim e^{-\gamma N}, \quad N \gg 1,$$

for some constant $\gamma > 0$ [41, 46, 48]. Next, since the collocation gPC is a spectrally accurate approximation of the polynomial projection in L^2 [18], then $\|f - g\|_2 \lesssim e^{-\gamma N}$ as well for $N \gg 1$. Finally, since $\|f - \pi_N(f)\|_\infty$ does not grow exponentially [18], Theorem 2.2 applies. \square

Two particularly important cases of this theorem are when ρ is the Lebesgue measure, associated with the Legendre polynomials ($\beta_{1,2} = 0$) and the measure associated with the Chebyshev polynomials ($\beta_{1,2} = -1/2$). By Theorem 4.3, the convergence of the Wasserstein metric stands in *sharp contrast* to that of the PDFs, i.e., of $\|p_\mu - p_\nu\|_{L^q}$. As previously noted, the convergence of the PDFs for the gPC method has not been proved, and might not be obtained at all for moderate values of N [13]. It remains an open question whether Theorem 4.3 can be extended to measures with an unbounded support, such as the normal and the exponential distributions. Such a generalization might require a generalization of Theorem 2.2 to unbounded domains. We further note that Theorem 4.3 can be extended to measures ρ' that are bounded from above by ρ , see [14] for details.

Numerical example. We approximate the same function f , as defined in (4.1), and approximate it using polynomial interpolation at Gauss-Legendre quadrature points. Since f is analytic, Theorem 4.3 guarantees that the gPC-based ν converges exponentially in N to that of μ , see Figure 4.1(b). The convergence of the respective PDFs, on the other hand, is polynomial at best (see Figure 4.1(c)). Quantitatively, the W_1 error decreases by 8 orders of magnitude between $N = 4$ and $N = 120$, whereas the L^1 distance between the PDFs decreases by only 4 orders of magnitude.

4.3. Comparison to the histogram method. This paper, as noted, is motivated by the following class of algorithms: to approximately characterize $\mu = f_*\rho$, first approximate f by g , and then approximate μ by $\nu = g_*\rho$. How does this approach compare with more standard statistical methods?

We focus on one common nonparametric statistical density estimation method, the histogram method; Given i.i.d. samples from μ , denoted by $f(\alpha_1) = y_1, \dots, f(\alpha_N) = y_N$, and a partition of the range of $f(\alpha)$ into L disjoint intervals (bins) $\{B_\ell\}_{\ell=1}^L$, the histogram estimator of the PDF is

$$p_{\text{hist}}(y) := \frac{1}{N} \sum_{\ell=1}^L (\# \text{ of samples for which } y_j \in B_\ell) \cdot \mathbb{1}_{B_\ell}(y),$$

where $\mathbb{1}_{B_\ell}$ is the characteristic function of bin B_ℓ [47]. The histogram method is intuitive and easy to implement. What is then the advantage of approximation-based UQ methods? In Sec. 4.4, using results by Bobkov and Ledoux [5], we prove the following corollary

COROLLARY 4.1. *Under the conditions of Theorem 4.1, the d -dimensional, m -th order spline-based estimator of μ outperforms the histogram method on average in the W_p sense when $d < 2(m + 1)$.*

The average in this corollary refers to all i.i.d. realizations of y_1, \dots, y_N from μ . This corollary is an example of the so-called ‘‘curse of dimensionality’’. To maintain a constant resolution and accuracy, the amount of data points (and hence the computational complexity) needs to increase exponentially with the dimension. Hence, above a certain dimension, it is preferable to ignore the underlying structure (i.e., the approximation of f by g) and to consider only the empirical distribution of the i.i.d. samples $\{f(\alpha_j)\}_{j=1}^N$.

4.4. Proof of Corollary 4.1. Given N i.i.d. from μ , denoted by y_1, \dots, y_N , define the empirical distribution as

$$\mu_{\text{emp}} := \frac{1}{N} \sum_{j=1}^N \delta_{y_j},$$

where δ_y is the Dirac delta distribution centered at the point $y \in \mathbb{R}$. Under certain broad assumptions (see [5] for details), $\mathbb{E}W_p(\mu, \mu_{\text{emp}}) \lesssim N^{-1/2}$, where the expectancy in these bounds is over all realizations of y_1, \dots, y_N with respect to the measure μ [5].

By the triangle inequality and linearity of expectation,

$$\mathbb{E}W_p(\mu, \mu_{\text{hist}}) \leq \mathbb{E}[W_p(\mu, \mu_{\text{emp}}) + W_p(\mu_{\text{emp}}, \mu_{\text{hist}})] = \mathbb{E}W_p(\mu, \mu_{\text{emp}}) + \mathbb{E}W_p(\mu_{\text{emp}}, \mu_{\text{hist}}),$$

where $d\mu_{\text{hist}}(y) = p_{\text{hist}}(y)dy$ is the measure defined by the histogram estimator. It is therefore sufficient to show that $\mathbb{E}W_p(\mu_{\text{emp}}, \mu_{\text{hist}}) \lesssim N^{-(1+1/p)}$ for any $p \geq 1$. We will prove a slightly stronger claim – that $W_p(\mu_{\text{emp}}, \mu_{\text{hist}}) \lesssim N^{-(1+1/p)}$ for every set of numbers y_1, \dots, y_N .

Let $\{B_\ell\}_{\ell=1}^L$ be the bins of the histogram estimator and let $\mu_{\text{emp},\ell}$ and $\mu_{\text{hist},\ell}$ be the restriction of the measures μ_{emp} and μ_{hist} to B_ℓ , respectively, for every $1 \leq \ell \leq L$. By definition, there are exactly $N \cdot \mu_{\text{hist}}(B_\ell)$ samples that fall into B_ℓ , and so $\mu_{\text{hist},\ell}(B_\ell) = \mu_{\text{emp},\ell}(B_\ell)$. Hence, the two measures $\mu_{\text{emp},\ell}$ and $\mu_{\text{hist},\ell}$ are comparable in the Wasserstein metric and we can write that

$$W_p^p(\mu_{\text{emp}}, \mu_{\text{hist}}) \leq \sum_{\ell=1}^L W_p^p(\mu_{\text{emp},\ell}, \mu_{\text{hist},\ell}).$$

Since $\mu_{\text{hist},\ell}$ is uniform on B_ℓ for any ℓ , the Wasserstein distance is the greatest if all of the samples in B_ℓ are located on the extreme edge of the bin, i.e., if $y_j \in B_\ell$ then $y_j = a_\ell$, where we denote $B_\ell = [a_\ell, b_\ell]$. Hence, for every $1 \leq \ell \leq L$,

$$\begin{aligned} W_p^p(\mu_{\text{emp},\ell}, \mu_{\text{hist},\ell}) &\leq \mu_{\text{emp},\ell}(B_\ell) \int_{a_\ell}^{b_\ell} (y - a_\ell)^p dy \\ &= \frac{\mu_{\text{emp},\ell}(B_\ell)}{p+1} (b_\ell - a_\ell)^{p+1}, \end{aligned}$$

and so

$$\begin{aligned} W_p^p(\mu_{\text{emp}}, \mu_{\text{hist}}) &\leq \sum_{\ell=1}^L \frac{\mu_{\text{emp},\ell}(B_\ell)}{p+1} (b_\ell - a_\ell)^{p+1} \\ &\lesssim N^{-(p+1)} \sum_{\ell=1}^L \mu_{\text{emp},\ell}(B_\ell) \\ &= N^{-(p+1)} \sum_{\ell=1}^L \mu_{\text{emp}}(B_\ell) = N^{-(p+1)}, \end{aligned}$$

where the second inequality is due to the partition, in which $(b_\ell - a_\ell) \sim N^{-1}$, and the last equality holds since μ_{emp} is a probability measure and since $\{B_\ell\}_{\ell=1}^L$ is a partition of its support.

Acknowledgments. The author would like to thank S. Steinerberger for many useful comments and advice. This research was partially carried out during a stay of the author as a guest of R.R. Coifman and the Department of Mathematics at Yale University, whose hospitality is gratefully acknowledged.

REFERENCES

- [1] M.J. Ablowitz and T.P. Horikis, *Interacting nonlinear wave envelopes and rogue wave formation in deep water*, Phys. Fluids, **27**:012107, 2015. 1.3
- [2] R.A. Adams and J.J.F. Fournier, *Sobolev Spaces*, Academic Press, New York, 2003. 2.2
- [3] M. Arjovsky, S. Chintala, and L. Bottou, *Wasserstein generative adversarial networks*, Proc. Int. Conf. Mach. Learn., **70**:214–223, 2017. 1.4
- [4] J. Beck, L. Tamellini, and R. Tempone, *IGA-based multi-index stochastic collocation for random PDEs on arbitrary domains*, Comput. Meth. Appl. Mech. Engrg., **351**:330–350, 2019. 4.1
- [5] S. Bobkov and M. Ledoux, *One-Dimensional Empirical Measures, Order Statistics and Kantorovich Transport Distance*, Memoirs of the Amer. Math. Soc., 2019. 4.3, 4.4
- [6] S. Boussaid, *work in progress*. 2.1
- [7] J.A. Carrillo, M. Difrancesco, A. Figalli, T. Laurent, and D. Slepcev, *Global-in-time weak measure solutions, finite-time aggregation and confinement for nonlocal interaction equations*, Duke Math. J., **156**:229–271, 2011. 1.4
- [8] Q.Y. Chen, D. Gottlieb, and J.S. Hesthaven, *Uncertainty analysis for the steady-state flows in a dual throat nozzle*, J. Comput. Phys., **204**:378–398, 2005. 1.3
- [9] I. Colombo, F. Nobile, G. Porta, A. Scotti, and L. Tamellini, *Uncertainty quantification of geochemical and mechanical compaction in layered sedimentary basins*, Comput. Methods Appl. Mech. Engrg., **328**:122–146, 2018. 1.3
- [10] P.J. Davis and P. Rabinowitz, *Numerical Integration*, Blaisdell, Waltham, MA, 1968. 4.2
- [11] C. De Boor, *A Practical Guide to Splines*, Springer-Verlag, New York, 1978. 4.1, 4.2, 4.1
- [12] L. Devroye and L. Györfi, *Nonparametric Density Estimation - The L_1 View*, Wiley, New York, 1985. 1.3
- [13] A. Ditekowski, G. Fibich, and A. Sagiv, *Density estimation in uncertainty propagation problems using a surrogate model*, SIAM/ASA J. Uncertain. Quantif., **8**(1):261–300, 2020. 1.2, 1.3, 1.1, 1.3, 3.4, 4, 4.1, 4.2, 4.2
- [14] A. Ditekowski and R. Kats, *On spectral approximations with nonstandard weight functions and their implementations to generalized chaos expansions*, J. Sci. Comput., **79**:1981–2005, 2019. 4.2
- [15] B. Ganapathysubramanian and N. Zabaras, *Sparse grid collocation schemes for stochastic natural convection problems*, J. Comput. Phys., **225**:652–685, 2007. 1.3
- [16] A.L. Gibbs and F.E. Su, *On choosing and bounding probability metrics*, Int. Stats. Rev., **70**:419–435, 2002. 1.3, 4.1, 3
- [17] C.R. Givens and R.M. Michael, *A class of Wasserstein metrics for probability distributions*, Michigan Math. J., **31**:231–240, 1984. 1.4
- [18] D. Gottlieb and S.A. Orszag, *Numerical Analysis of Spectral Methods: Theory and Applications*, SIAM, Philadelphia, PA, USA, 1979. 4.2, 4.2
- [19] Y. van Halder, B. Sanderse, and B. Koren, *An adaptive minimum spanning tree multi-element method for uncertainty quantification of smooth and discontinuous responses*, SIAM J. Sci. Comput., **41**(6):A3624A3648, 2019. 1.3, 4.1
- [20] C.A. Hall and W.W. Meyer, *Optimal error bounds for cubic spline interpolation*, J. Approx. Theory, **16**:105–122, 1976. 4.2
- [21] O.P. Le Maître, L. Mathelin, O.M. Knio, and M.Y. Hussaini, *Asynchronous time integration for polynomial chaos expansion of uncertain periodic dynamics*, Discrete Contin. Dyn. Syst., **28**:199–226, 2010. 1.3
- [22] E.H. Leeb and M. Loss, *Analysis*, Graduate Studies in Mathematics, Amer. Math. Soc., Providence, RI, 14, 2001. 2.2
- [23] G. Loeper, *Uniqueness of the solution of the Vlasov-Poisson system with bounded density*, J. Math. Pures Appl., **86**:68–79, 2006. 2.2, 2.3
- [24] R.J. McCann, *A convexity principle for interacting gases*, Adv. Math., **128**:153–179, 1997. 1.4
- [25] K. Ni, X. Bresson, T. Chan, and S. Esedoglu, *Local histogram based segmentation using the Wasserstein Distance*, Int. J. Comp. Vis., **84**:97–111, 2009. 1.4
- [26] G. Patwardhan, X. Gao, A. Sagiv, A. Dutt, J. Ginsberg, A. Ditekowski, G. Fibich, and A.L. Gaeta, *Loss of polarization in elliptically polarized collapsing beams*, Phys. Rev. A, **99**:033824, 2019.

- 1.3
- [27] R. Peyre, *Comparison between W^2 distance and \dot{H}^{-1} , and localization of Wasserstein distance*, ESAIM Control Optim. Calc. Var., **24:1489–1501**, 2018. 2.2, 2.3
 - [28] G. Pólya and G. Szegő, *Isoperimetric Inequalities in Mathematical Physics*, Princeton University, NJ, 1951. 2.2
 - [29] P.M. Prenter, *Splines and the Variational Method*, Wiley, New York, 2008. 4.1
 - [30] J.R. Rice, *Multivariate piecewise polynomial approximation*, In D.G. Handscomb (eds.), *Multivariate Approximation*, Academic Press, New York, 261–277, 1978. 4.1
 - [31] Y. Rubner, C. Tomasi, and L.J. Guibas, *The earth mover’s distance as a metric for image retrieval*, Int. J. Comp. Vis., **40, 99–121**, 2000. 1.4
 - [32] A. Sagiv, A. Ditzkowski, and G. Fibich, *Loss of phase and universality of stochastic interactions between laser beams*, Opt. Exp., **25:24387–24399**, 2017. 1.3
 - [33] T. Salvemini, *Sul calcolo degli indici di concordanza tra due caratteri quantitativi*, Atti della I Riunione della Soc. Ital. di Statistica, Roma, 1943. 1.2
 - [34] F. Santambrogio, *Optimal Transport for Applied Mathematicians. Calculus of Variations, PDEs, and Modeling, Progress in Nonlinear Differential Equations and their Applications*, Birkhäuser, New York, 2015. 1.4, 1.4
 - [35] M.H. Schultz, *L^∞ -multivariate approximation theory*, SIAM J. Numer. Anal., **6:161–183**, 1969. 4.1, 4.1
 - [36] S. Solem, *Convergence rates of the front tracking method for conservation laws in the Wasserstein distances*, SIAM J. Numer. Anal., **56:3648–3666**, 2018. 1.4
 - [37] S. Steinerberger, *Wasserstein distance, Fourier series and applications*, arXiv preprint, arxiv:1803.08011, 2018. 2.2
 - [38] B. Sudret and A. Der Kiureghian, *Stochastic finite element methods and reliability: A state-of-the-art report*, Department of Civil and Environmental Engineering, University of California Berkeley, 2000. 1.3
 - [39] G. Szego, *Orthogonal Polynomials*, Colloquium Publications, Amer. Math. Soc., New York, **23**, 1939. 4.2
 - [40] E. Tadmor, *Local error estimates for discontinuous solutions of nonlinear hyperbolic equations*, SIAM J. Numer. Anal., **28:891–906**, 1991. 1.4
 - [41] L.N. Trefethen, *Approximation Theory and Approximation Practice*, SIAM, Philadelphia, PA, 2013. 4.2
 - [42] A.B. Tsybakov, *Introduction to Nonparametric Estimation*, Springer, New York, 2009. 1.3
 - [43] S. Ullmann, *POD-Galerkin modeling for incompressible flows with stochastic boundary conditions*, M.Sc. Dissertation, Technical University of Darmstadt, 2014. 1.3
 - [44] S.S. Vallender, *Calculation of the Wasserstein distance between probability distributions on the line*, SIAM Theory Probab. Appl., **18:784–786**, 1974. 1.2
 - [45] C. Villani, *Topics in Optimal Transportation*, Amer. Math. Soc., 2003. 1.4, 1.4, 3.3
 - [46] H. Wang and S. Xiang, *On the convergence rates of Legendre approximation*, Math. Comp., **81:861–877**, 2012. 4.2
 - [47] L. Wasserman, *All of Statistics: A Concise Course in Statistical Inference*, Springer Science & Business Media, New York, 2004. 4.3
 - [48] D. Xiu, *Numerical Methods for Stochastic Computations: A Spectral Method Approach*, Princeton University, NJ, 2010. 1.3, 4.2, 4.2
 - [49] D. Xiu and J.S. Hesthaven, *High-order collocation methods for differential equations with random inputs*, SIAM J. Sci. Comput., **27:1118–1139**, 2005. 1.3, 4.2
 - [50] D. Xiu and G.E. Karniadakis, *The Wiener–Askey polynomial chaos for stochastic differential equations*, SIAM J. Sci. Comput., **24:619–644**, 2002. 1.3