

APPROXIMATION THEORY, COMPUTING, AND DEEP LEARNING ON THE WASSERSTEIN SPACE

MASSIMO FORNASIER, PASCAL HEID, AND GIACOMO ENRICO SODINI

ABSTRACT. The challenge of approximating functions in infinite-dimensional spaces from finite samples is widely regarded as formidable. In this study, we delve into the challenging problem of the numerical approximation of Sobolev-smooth functions defined on probability spaces. Our particular focus centers on the Wasserstein distance function, which serves as a relevant example. In contrast to the existing body of literature focused on approximating efficiently pointwise evaluations, we chart a new course to define functional approximants by adopting three machine learning-based approaches:

1. Solving a finite number of optimal transport problems and computing the corresponding Wasserstein potentials.
2. Employing empirical risk minimization with Tikhonov regularization in Wasserstein Sobolev spaces.
3. Addressing the problem through the saddle point formulation that characterizes the weak form of the Tikhonov functional's Euler-Lagrange equation.

As a theoretical contribution, we furnish explicit and quantitative bounds on generalization errors for each of these solutions. In the proofs, we leverage the theory of metric Sobolev spaces and we combine it with techniques of optimal transport, variational calculus, and large deviation bounds. In our numerical implementation, we harness appropriately designed neural networks to serve as basis functions. These networks undergo training using diverse methodologies. This approach allows us to obtain approximating functions that can be rapidly evaluated after training. Consequently, our constructive solutions significantly enhance at equal accuracy the evaluation speed, surpassing that of state-of-the-art methods by several orders of magnitude.

CONTENTS

1	Introduction	2
2	Wasserstein Sobolev spaces	11
2.1	The relaxation approach to Metric Sobolev spaces	11
2.2	Wasserstein spaces	12
2.3	Density of subalgebras of Lipschitz functions in Wasserstein Sobolev spaces	13
3	Approximability of Wasserstein distances by cylinder maps	15
3.1	Distribution independent approximation	15
3.2	Random subcoverings	18
4	Consistency of empirical risk minimization	25
4.1	Gamma convergence with approximating measures	25
4.2	Regularized least squares approximation of bounded functions	29
5	Solving Euler–Lagrange equations	36
5.1	Corresponding saddle point problem	38
6	Numerical experiments	40
6.1	Wasserstein distance approximation for a discrete base set	41
6.2	Approximation by a deep neural network	47

2020 *Mathematics Subject Classification*. Primary: 49Q22, 33F05; Secondary: 46E36, 28A33, 68T07.

Key words and phrases. Wasserstein Sobolev spaces, Approximation theory, Kantorovich-Wasserstein distance, Optimal Transport, Empirical risk minimization, Tikhonov regularization, Generalization error, Deep learning.

6.3 Adversarial deep neural network approach for the solution of the saddle-point problem	53
References	55

1. INTRODUCTION

In this work we are concerned with the efficient numerical approximation of Sobolev-smooth functions defined on spaces of probability measures from the information obtained by a finite number of point evaluations. Such approximation problems are considered already very challenging for functions defined on high dimensional Euclidean spaces, but they become particularly intriguing and formidable as the domain is a metric space of infinite dimensional nature. As an inspiring and motivating example, we focus in particular - although not solely, see Section 4 and Section 5 - on the study of the approximation of the Wasserstein distance function $\mu \mapsto W_p(\mu, \vartheta)$, where $\vartheta \in \mathcal{P}(K)$ is a given reference measure on a compact set $K \subset \mathbb{R}^d$, both from a theoretical and computational point of view. We recall that the Wasserstein distance arises as the solution of an optimal transport problem; cf. Section 2.2 below. Our theoretical approximation bounds are largely based on the theory of Wasserstein Sobolev spaces, especially on their Hilbertian structure, the Cheeger energy, and the algebra of cylinder functions; those notions are recalled in Section 2, see also [37, 86]. Concerning the computational framework, we make use of the fact that the algebra of cylinder functions is dense in the Wasserstein Sobolev space and that every cylinder function can be approximately realized by deep neural networks.

Hence, from a foundational point of view, with this paper we contribute to pioneer the connection of metric measure space theory and metric Sobolev spaces [4, 20, 24, 84] with numerical computations and machine learning. As we tread this novel route, we draw some conceptual inspiration from previous works such as [5, 93, 52] that enable the spectral embedding of RCD spaces into L^2 spaces. Especially in cases where the spectrum of the Laplacian is discrete, as in the compact case, it offers an intriguing tool for the Euclidean embedding of metric space data points. From a computational point of view, the strength of our approach is the fast evaluation of Sobolev-smooth functions such as the Wasserstein distance $\mu \rightarrow W_p(\mu, \vartheta)$ once a deep neural network is setup and suitably trained. In particular, in perspective of the evaluation time, our approach outperforms the commonly employed methods for the computation of the Wasserstein distance.

Let us now recall the relevance of the Wasserstein distance and the commonly used approaches for its efficient pointwise evaluation. The Wasserstein distance is an important tool to compare (probability) measures. Besides its crucial role in the theory of optimal transport [79, 97], it has had by now far reaching uses in numerical applications [76]. One of the first data-driven applications of a discrete Wasserstein metric was in image retrieval and can be found in [78]. Since then, the Wasserstein metric has been successfully applied in various areas such as image processing, computer vision, statistics, and machine learning; we refer to [61] for a comprehensive overview of specific applications and further references. The investigation of computational methods for the Wasserstein distance and similar functions acting on spaces of probability measures has gained significant prominence in recent years due to its crucial role in various practical applications. This field of research has witnessed substantial progress. Regrettably, in our overview below, we may have inadvertently omitted some recent works in this area; we apologize to those whose contributions may not have been included in this discussion. In the following, we treat the computational schemes for the three possible cases (a) both measures are continuous, (b) both measures are discrete, i.e., a sum of finitely many weighted Dirac measures, and (c) one is continuous and the other discrete, separately.

We start our overview by presenting some of the existing procedures for the numerical solution of the Wasserstein distance in the context of two discrete measures; this is certainly the best studied and most elaborated among the three cases (a)–(c) outlined above. First of all, we note that in the discrete setting the Kantorovich problem is a classical, finite dimensional, linear program; i.e., the problem amounts to the optimization of an objective function that is linear and whose constraints are linear as well. Indeed, the Kantorovich problem is a minimum cost network flow problem. Consequently, in the given setting, the Wasserstein distance can be computed by common algorithmic tools from linear programming. A very well written overview of applicable algorithms is given in the book of Peyré and Cuturi, [76], or the thesis [82]. To mention only a few of them, we point to the Hungarian method introduced by Kuhn, [64], or the Auction method, which was originally proposed by Bertsekas, [17], further improved in [19], and applied to the transportation problem in [18]; we further refer to [22]. Even though those algorithms allow to compute an approximation with an arbitrary accuracy, a common disadvantage of those algorithms stemming from the linear programming is that they are computationally very expensive; see, e.g., [43, 92, 75]. In particular, they have cubical complexity (with respect to the number of atoms of the discrete measures). The computational cost can be drastically improved by solving a flow problem on a suitable graph, see [67] and the more recent work [11], which, however, only applies to the 1-Wasserstein distance, but not to a general p -Wasserstein distance. A prominent approach to relieve the computational cost is to add an entropic regularization penalty term to the optimal transport problem. We refer to the book of Peyré and Cuturi, [76], for a gentle treatment of this topic. We stress that the solution of the regularized problem converges to the optimal transport plan as the regularization parameter goes to zero. The most famous approach to solve the regularized problem is given by the Sinkhorn algorithm, which applies a simple iteration procedure, see, e.g., [76, Sec. 4.2]. We note that these iterations are based on matrix-vector products, and thus are suitable for GPU computations. Furthermore, this scheme was formerly known as the iterative proportional fitting procedure (IPFP), see, e.g., [32]. However, as the convergence proof is credited to R. Sinkhorn, [85], the iteration method is nowadays known as the Sinkhorn algorithm. It was shown in [38] that the convergence of the Sinkhorn algorithm is indeed linear. The Sinkhorn algorithm has regained more attention thanks to the work [28], which highlighted that this scheme has several favourable computational properties. Since then, the algorithm has witnessed many modifications and extensions. For instance, the numerical instability with respect to a small regularization parameter was alleviated by Schmitzer, [81]. In [16], the authors further improved Sinkhorn’s algorithm by exploiting that the regularized problem corresponds to the Kullback–Leibler Bregman divergence projection, which allows to apply a Bregman–Dykstra iteration scheme. We further remark that for the approximation of the solution of the *unregularized* problem one may apply a so-called proximal point algorithm for the Kullback–Leibler divergence, cf. [76, Rem. 4.9]. Finally, for a generalized Sinkhorn algorithm that can be applied to larger class of convex optimization problems we refer to [76, Sec. 4.6].

A more recent approach to compute the 2-Wasserstein distance is given by the linear optimal transport (LOT) framework, which was originally introduced in [98] in the context of image datasets; see also the follow-up works [12, 62]. This approach employs a meaningful linearization of the Wasserstein metric measure space by a projection onto the tangent space at a given reference measure. Then, to compute the LOT distance between two measures, which in particular is an Euclidean distance on the tangent space, one first needs to compute the optimal transport plan of each of those measures to the fixed reference measure. We also refer to the closely related work [83]. Moreover, in [27], an Euclidean embedding, for which the Euclidean distance again approximates the Wasserstein distance, is learnt with a neural network. For further applications of Wasserstein embeddings we refer, for instance, to [60, 54]; certainly, there are (many) more works that apply the Wasserstein embedding methodology.

At last in the context of the computation of the Wasserstein distance between two discrete measures, we shall point to the Wasserstein generative adversarial networks (WGANs) as introduced in [7]. In that framework, the Wasserstein distance is estimated by computing the Kantorovich dual for a discriminator network. More precisely, the set of admissible potentials in the dual formulation is restricted to the realization of a given neural network, which is the discriminator in the setting of WGAN. Subsequently, the neural network is trained to maximize the expected value in the dual formulation; i.e., the discriminator is trained to approximate the Kantorovich potential. This approach has indeed been very successful. We further refer to the improved WGAN [47], and to [68], where two-step method for the computation of the Wasserstein distance in WGANs is introduced. In [73] a comparison of different schemes studied in the WGAN literature is presented, and [89] provides theoretical insights of WGANs.

Next, we address the semi-discrete setting, which, however, is only discussed superficially; we point the interested reader to [76, 79] and the references given therein for more details. A common approach to solve the semi-discrete optimization problem numerically is to consider the dual problem and split the integral over the underlying domain into so-called Laguerre cells. In the special case $p = 2$, i.e., the 2-Wasserstein distance, the Laguerre cells are known as *power cells*. In that case, the cells are polyhedral, and can be computed efficiently using computational geometry algorithms, see, e.g., [9]. Then, in order to solve the optimization problem, one may apply, combined with a scheme from computational geometry, gradient algorithms, which include Netwon methods; we refer, for instance, to the works [74, 65, 57, 72, 58, 10].

Lastly, we consider the *continuous* case, specifically the computation of the Wasserstein distance between (probability) distributions. For that purpose, we mainly summarize the schemes presented in [79, Sec. 6] in a condensed manner, and recall some of the reference stated therein. The most prominent method in the continuous case is based on the Benamou–Brenier formulation, in which the optimal transport problem is transformed into a convex optimization problem with linear constraints; see, e.g., the monographs [79, 1], or the original report [14]. In order to solve this convex optimization problem numerically, the authors of [14] employed an augmented Lagrangian scheme combined with Uzawa’s gradient method. It can be verified that this procedure converges. For instance, this was shown in the works of Benamou, Brenier, and Guittet, [15], and Guittet, [46]; we refer to [79, Sec. 6.1] and the references therein for further details and other application areas of this approach. Further numerical schemes in the context of continuous measures include the algorithm due to Angenent, Haker, and Tannenbaum, cf. [6] and [48], as well as computational procedures for the Monge–Ampère problem. This equation was solved in [69] by an iterative nonlinear solver based on Newton’s method, whereas the authors of [23] employed a gradient descent solution to the Monge–Ampère problem. We shall further mention the manuscript [25], in which the authors proved a sample complexity bound and studied the performance of the Sinkhorn divergence estimator, which was introduced in [77].

Finally, we want to point to the stochastic optimization approach for large-scale optimal transport problems as introduced in [39], which can be applied to the continuous, discrete, as well as the semi-discrete cases. Especially in the continuous setting, the authors of [39] proposed a novel method that exploits reproducing kernel Hilbert spaces, and proved the convergence under suitable assumptions.

In contrast to the existing body of literature, we move from the commonly addressed problem of approximating the Wasserstein distance pointwise to the more abstract one of approximating it as a function. In particular, our work takes a fresh direction by adopting a machine learning approach for computing the Wasserstein distance as a function on datasets. To our knowledge, this is the first work in this direction. Specifically, our objective is to compute this distance

by training a straightforward approximation function using a finite training set of Wasserstein distance evaluations from a dataset, while ensuring a minimal relative test error across the remaining dataset. Additionally, we aim for the approximating function to possess a simplicity that allows for rapid numerical evaluations after training, which could improve of several orders of magnitude the speed of evaluation with respect to state-of-the-art methods. As we elaborate on later, our approach extends beyond the specific case of the Wasserstein distance to encompass more general Wasserstein Sobolev functions.

Throughout the paper, we delve into the examination of three closely interconnected machine learning approaches for approximating Wasserstein Sobolev functions by

1. solving a finite number of optimal transport problems and computing the corresponding Wasserstein potentials, cf. Section 3 and Section 6.1 (this approach is limited to the approximation of the Wasserstein distance);
2. empirical risk minimization using a Tikhonov regularization in Wasserstein Sobolev spaces, cf. Section 4;
3. solving the saddle point problem that describes the Euler-Lagrange equation in weak form of the Tikhonov functional, cf. Section 5.

In all these solutions we employ suitably defined neural networks to implement the basis functions. In the following we consider both generic non-negative and finite Borel measures \mathbf{m} over the space $\mathcal{P}(K)$ and probability measures $\mathbf{p} \in \mathcal{P}(\mathcal{P}(K))$ in order to describe data distributions; to keep things as simple as possible, let assume the domain to be a compact subset $K \subseteq \mathbb{R}^d$. The choice between these two options depends on whether the results are deterministic or are based on randomization of the input. The starting point for our analysis are the results of [37, 86] concerning the approximability of the Wasserstein distance and other Wasserstein Sobolev functions by so called *cylinder functions*. Let us also mention that, even if we limit our analysis to the space of probability measures, many of the techniques developed in [37, 86] can also be applied to the space of non-negative and finite measures endowed with the Hellinger-Kantorovich distance ([87, 66]). For that reason we expect that the analysis developed herein should be robust enough so that it can be further generalized in this direction.

As a starting and relevant example, the function that we would like to approximate is given by

$$F_{\vartheta}^{W_2}(\mu) := W_2(\mu, \vartheta), \quad \mu \in \mathcal{P}(K), \quad (1.1)$$

where $\vartheta \in \mathcal{P}(K)$ is a fixed reference probability measure and W_2 is the 2-Wasserstein distance between probability measures (while in the work we often consider a parameter p possibly different from 2 for the Wasserstein distance, we prefer to consider the simpler case $p = 2$ at this introductory level), arising as the solution of the optimal transport problem

$$W_2(\mu, \vartheta) := \left(\inf \left\{ \int_{K \times K} |x - y|^2 d\gamma(x, y) : \gamma \in \mathcal{P}(K \times K) \text{ with marginals } \mu \text{ and } \vartheta \right\} \right)^{1/2}.$$

One of the main results of [37] states that, given any non-negative and finite Borel measure \mathbf{m} on $\mathcal{P}(K)$, it is possible to find a sequence $F_n : \mathcal{P}(K) \rightarrow \mathbb{R}$ of functions having the form

$$F_n(\mu) := \psi_n \left(\int_K \phi_1^n d\mu, \int_K \phi_2^n d\mu, \dots, \int_K \phi_{N_n}^n d\mu \right), \quad \mu \in \mathcal{P}(K) \quad (1.2)$$

such that

$$F_n \rightarrow W_2(\cdot, \vartheta) \text{ in } L^2(\mathcal{P}(K), \mathbf{m}), \quad (1.3)$$

where ψ_n and ϕ_i^n are smooth functions in appropriate Euclidean spaces. A function of the form (1.2) is called *cylinder function*, see also (2.9). For later use, we denote by $\mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$ the set of all cylinder functions. We anticipate now that, as long as the function ψ_n is realized or well-approximated by a neural network, the cylinder function F_n can itself be considered a neural network over $\mathcal{P}(K)$ with first-layer weights ϕ_i^n used to integrate the input μ . The efficient

approximation of high-dimensional functions, such as ψ_n , by neural networks is underlying our numerical results (cf. Section 6), but it is not central or exhaustively studied in our theoretical analysis. Instead, we do refer to the large body of more specific literature dedicated to the approximation theory by neural networks [59, 29, 51, 50, 21, 34, 55, 49, 33], which is far from an exhaustive list. To provide a more explicit disclaimer regarding the use of neural networks, it is important to emphasize that we do not specifically target ensuring global optimization when dealing with neural networks. Given the inherent non-convex nature of the problem, addressing generalization errors in neural network optimization is widely recognized as one of the most challenging aspects of machine learning. In this paper, our focus does not extend to tackling this intricate issue. Instead, in our numerical experiments, we make the simplifying assumption that we are aiming at approximating minimization, and our primary goal is to achieve values of the objective functions that are sufficiently small.

The approximation result by cylinder functions has deep consequences in terms of the structure of the so called *metric Sobolev space* on $(\mathcal{P}(K), W_2, \mathfrak{m})$ that we briefly discuss in Section 2 (and we refer the reader to [37, 86] for a more extended analysis), but here we are more interested in the practical consequences of this approximation property. In particular, while computing numerically the Wasserstein distance between two measures may be a challenging and heavy task, the computation of the above functions F_n consists simply in the calculation of a finite number of integrals and then in the evaluation of a function of the resulting values. The simplicity of the latter procedure is thus the main reason to study how to provide an explicit and numerically efficient construction of the functions F_n , also because the convergence proof elaborated in [37] and further developed in [86] is not fully constructive.

This is the first problem that we address in Section 3. In particular, we show (see Proposition 3.4) that, up to knowing a *countable* number of functions (called *Kantorovich potentials*), it is possible to build maps ψ_n and ϕ_i^n satisfying the convergence (1.3); the main idea comes from the celebrated Kantorovich duality theorem, see, e.g., [3, Theorem 6.1.4], stating that, for every pair of sufficiently well-behaved probability measures $\mu, \vartheta \in \mathcal{P}(K)$, there exists a pair of potentials $\varphi^*, \psi^* : K \rightarrow \overline{\mathbb{R}}$ such that

$$W_2(\mu, \vartheta) = \left(\int_K \varphi^* d\mu + \int_K \psi^* d\vartheta \right)^{1/2} = \sup_{\substack{\varphi(x) + \psi(y) \leq |x-y|^2, \\ \varphi, \psi \in C_b(K)}} \left(\int_K \varphi d\mu + \int_K \psi d\vartheta \right)^{1/2}.$$

In general, however, the pair (φ^*, ψ^*) depends on ϑ and μ so that it is not possible to simply consider the function $F : \mathcal{P}(K) \rightarrow \mathbb{R}$ built as

$$F(\nu) := \left(\int_K \varphi^* d\nu + \int_K \psi^* d\vartheta \right)^{1/2}, \quad \nu \in \mathcal{P}(K),$$

as a cylinder function approximating the Wasserstein distance between ν and ϑ . This first issue can be fixed by considering a dense and countable subset of measures $(\sigma^h)_h \subset \mathcal{P}(K)$ and the supremum over their corresponding Kantorovich potentials. The second issue concerns the regularity of such potentials, which, in general, are neither smooth nor bounded; it is thus useful to consider specific truncations and regularization arguments in order to produce a suitable sequence of potentials.

All in all, the first result we obtain is Proposition 3.4, which we report here in a simplified version for the sake of this introduction.

Proposition 1.1. *Let $K \subset \mathbb{R}^d$ be a compact set and let $\vartheta \in \mathcal{P}(K)$. There exist a sequence of smooth functions $(\eta_n)_n \subset C_b^1(\mathbb{R}^n)$ and smoothed versions of Kantorovich potentials $(v_i^n)_n \subset C_b^1(K)$,*

$1 \leq i \leq n$, (built starting from a dense subset $(\sigma^h)_h \subset \mathcal{P}(K)$) such that the function

$$F_n(\mu) := \eta_n \left(\int_K v_i^n d\mu, \int_K v_2^n d\mu, \dots, \int_K v_n^n d\mu \right), \quad \mu \in \mathcal{P}(K),$$

converges pointwise monotonically from below to $W_2(\vartheta, \cdot)$ in $\mathcal{P}(K)$ as $n \rightarrow \infty$.

The pointwise convergence obtained above in turn implies

$$F_n \rightarrow W_2(\cdot, \vartheta) \text{ in } L^2(\mathcal{P}(K), \mathbf{m}) \text{ as } n \rightarrow \infty,$$

for any measure \mathbf{m} . Hence, it is a result that ensures approximability in $L^2(\mathcal{P}(K), \mathbf{m})$ in a universal manner. It is certainly a solid first step in the direction of the computability of the smooth approximation of the Wasserstein distance by cylinder functions, but it still suffers from two limitations: first of all it requires the knowledge of an infinite number of Kantorovich potentials and, secondly, it does not come with any quantitative convergence rate.

We obtain both these improvements by trading such a deterministic construction and its universality with a probabilistic approach that depends on the choice of the underlying measure \mathbf{m} ; namely, in Section 3.2, we introduce the notion of random subcovering of a metric measure space and we use it to quantify the convergence of (a randomly adapted version of) the functions F_n as above.

Let us first of all introduce the relevant definition: given a probability measure \mathbf{p} on a complete and separable metric space (\mathcal{S}, d) (as a concrete example we have the metric space $(\mathcal{P}(K), W_2)$ in mind), we call the quantity

$$p_{\varepsilon, n} := \mathbb{P} \left(X \in \cup_{i=1}^n B(X_i, \varepsilon) \right),$$

the (ε, n) -subcovering probability, where X, X_1, \dots, X_n are drawn i.i.d. according to \mathbf{p} , $\varepsilon \in (0, 1)$ and $n \in \mathbb{N}$. This is simply the probability that a (random) point $X \in \mathcal{S}$ belongs to at least one of the (random) ε -balls centered at points X_1, \dots, X_n and serves as a quantification of the concentration of the measure \mathbf{p} .

After providing a few results concerning the study of the quantity $p_{\varepsilon, n}$, see in particular Lemma 3.7 showing that $p_{\varepsilon, n} \rightarrow 1$ as $n \rightarrow +\infty$, the main result of Section 3.2 is the following ‘‘randomized and quantified’’ version of the above Proposition 1.1. Again, the version reported here is simplified for the sake of clarity, and we refer to Proposition 3.13 for the complete statement.

Proposition 1.2. *Let $N \in \mathbb{N}$ and $K \subset \mathbb{R}^d$ be a compact set. Consider μ_1, \dots, μ_N i.i.d. random variables on $\mathcal{P}(K)$ distributed according to $\mathbf{p} \in \mathcal{P}(\mathcal{P}(K))$ and let $\vartheta \in \mathcal{P}(K)$ be fixed. Let (v_1, \dots, v_N) be smoothed versions of Kantorovich potentials computed starting from the measures μ_1, \dots, μ_N and let $F_N : \mathcal{P}(K) \rightarrow \mathbb{R}$ be defined as*

$$F_N(\mu) := \eta_N \left(\int_K v_1 d\mu, \int_K v_2 d\mu, \dots, \int_K v_N d\mu \right), \quad \mu \in \mathcal{P}(K).$$

Then

$$\mathbb{E} \left[\int_{\mathcal{P}(K)} |F_N(\mu) - F_{\vartheta}^{W_2}(\mu)|^2 d\mathbf{p}(\mu) \right] \leq C [(1 - p_{\varepsilon, N}) + p_{\varepsilon, N}(N^{-2} + \varepsilon^2)],$$

where C is a constant depending only on K .

It is clear that the computation of the above function F_N is drastically simpler than the one of F_n as in Proposition 1.1, requiring the knowledge of only a finite number of Kantorovich potentials, which are drawn randomly; we can also see an explicit order of convergence in terms of the (ε, n) subcovering probabilities. In fact, thanks to the compactness of K , for $\varepsilon > 0$ fixed, we have that $p_{\varepsilon, N} \rightarrow 1$ exponentially fast as $N \rightarrow \infty$, see Lemma 3.7. We further mention that the estimate in Proposition 1.2 can be regarded as a *generalization error* as it quantifies in mean-squares the misfit of the approximant integrated over the entire measure \mathbf{p} .

This result is further studied in the specific context of a discrete base domain \mathfrak{D} ; i.e., the compact set K is replaced by a discrete set \mathfrak{D} . The restriction to a discrete domain allows for a simpler construction of the approximating sequence by cylinder functions, which still relies on pre-computed Kantorovich potentials; additionally, we show how to numerically and efficiently implement the procedure on a computer via neural networks (whose architecture is used also for training). In particular, for a fixed reference measure $\vartheta \in \mathcal{P}(\mathfrak{D})$ and a dense and countable subset $(\mu_k)_{k \in \mathbb{N}}$ of $\mathcal{P}(\mathfrak{D})$, denote by (φ_k, ψ_k) , $k \in \mathbb{N}$, the corresponding pairs of Kantorovich potentials; i.e., we have that

$$F_{\vartheta}^{W_2^2}(\mu_k) := W_2^2(\vartheta, \mu_k) = \int_{\mathfrak{D}} \varphi_k \, d\mu_k + \int_{\mathfrak{D}} \psi_k \, d\vartheta. \quad (1.4)$$

Based thereon, for $I \subset \mathbb{N}$, we define the cylinder function

$$G_{\vartheta}^I(\mu) := \max_{k \in I} \int_{\mathfrak{D}} \varphi_k \, d\mu + \int_{\mathfrak{D}} \psi_k \, d\vartheta. \quad (1.5)$$

We show (cf. Theorem 6.7) that, for any $\varepsilon > 0$ there exists a finite index set $I_{\varepsilon} \subset \mathbb{N}$ such that

$$\sup_{\mu \in \mathcal{P}(\mathfrak{D})} \left| F_{\vartheta}^{W_2^2}(\mu) - G_{\vartheta}^{I_{\varepsilon}}(\mu) \right| \leq \varepsilon. \quad (1.6)$$

Thereupon we verify that the cylinder function G_{ϑ}^I from (1.5) can be realized as a deep neural network. Indeed, since

$$\mu \mapsto \int_{\mathfrak{D}} \varphi_k \, d\mu + \int_{\mathfrak{D}} \psi_k \, d\vartheta$$

is an affine function, it can be represented in the finite dimensional setting by a (weight) matrix plus a (bias) vector. Moreover, it is known that the maxima function can be realized by a deep neural network, see, e.g., [13, 45, 53], and the composition of those functions can be obtained by the concatenation of the neural networks ([53, Prop. 2.14]).

The approach described by Proposition 1.2 does apply quite specifically to the approximation of the Wasserstein distance function, thanks to the Kantorovich duality. In order to allow for efficient approximations of more general functions, we study the construction of approximants by empirical risk minimization. In particular, in Section 4, our specific goal is to elucidate the conditions and the method through which we can achieve the best possible approximation of the Wasserstein distance function as in equation (1.1) and other Sobolev-smooth functions. We aim to do this within the context of the space $(\mathcal{P}(K), W_2, \mathbf{m}_n)$, where \mathbf{m}_n represents an empirical estimate of the measure \mathbf{m} . Furthermore, we introduce a degree of regularization using the pre-Cheeger energy as additional term to cope with noisy data. Ultimately, our objective is to showcase how this approximation converges robustly to the sought function within the space $(\mathcal{P}(K), W_2, \mathbf{m})$. We reiterate that this theory allows us to extend the scope of our results to a larger class of functions (beyond the Wasserstein distance function), which are elements of suitably defined Wasserstein Sobolev spaces. Before stating the precise result we need to briefly introduce the concept of (pre-)Cheeger energy and Sobolev norm in this context (see Section 2 for a more comprehensive explanation and references to the literature).

The fundamental idea is that to every cylinder function $F : \mathcal{P}(K) \rightarrow \mathbb{R}$, thus having the form

$$F(\mu) := \psi \left(\int_K \phi_1 \, d\mu, \dots, \int_K \phi_N \, d\mu \right), \quad \mu \in \mathcal{P}(K),$$

we can associate a notion of derivative given by

$$DF(\mu, x) := \sum_{n=1}^N \partial_n \psi(\mathbf{L}_{\phi}(\mu)) \nabla \phi_n(x), \quad (\mu, x) \in \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d, \quad (1.7)$$

where $L_\phi(\mu) := (\int_K \phi_1 d\mu, \dots, \int_K \phi_N d\mu)$. Besides being intuitively the correct object to be rightfully considered as the derivative of F , it can be seen that the $L^2(K, \mu)$ -norm of $DF(\mu, \cdot)$ corresponds to the so called *asymptotic Lipschitz constant* of F at μ (cf. Proposition 2.1), a purely metric object that generalizes the notion of derivative.

This leads to the definition of the *pre-Cheeger energy* of the function F given by

$$\mathfrak{pCE}_{2,\mathfrak{m}}(F) := \int_{\mathcal{P}(K)} \int_K |DF(\mu, x)|^2 d\mu(x) d\mathfrak{m}(\mu) = \int_{\mathcal{P}(K)} \|DF(\mu, \cdot)\|_{L^2(K, \mu)}^2 d\mathfrak{m}(\mu).$$

The completion of the space of cylinder functions with respect to the norm

$$|F|_{pH_{\mathfrak{m}}^{1,2}}^2 := \|F\|_{L^2(\mathcal{P}(K), \mathfrak{m})}^2 + \mathfrak{pCE}_{2,\mathfrak{m}}(F) \quad (1.8)$$

is the Sobolev space $H^{1,2}(\mathcal{P}(K), W_2, \mathfrak{m})$ with norm

$$|F|_{H_{\mathfrak{m}}^{1,2}}^2 := \|F\|_{L^2(\mathcal{P}(K), \mathfrak{m})}^2 + \mathfrak{CE}_{2,\mathfrak{m}}(F), \quad F \in H^{1,2}(\mathcal{P}(K), W_2, \mathfrak{m}),$$

where $\mathfrak{CE}_{2,\mathfrak{m}}$ is called *Cheeger energy* of F ; the specific definition of this notion is, for the moment, not relevant, but is introduced below in (2.4).

The main result of Section 4.1 is the following (we refer to Theorem 4.4 and Corollary 4.5 for the detailed and complete statements).

Theorem 1.3. *Let $(\mathfrak{m}_n)_n$ be a sequence of non-negative Borel measures on $\mathcal{P}(K)$ weakly converging to the measure \mathfrak{m} and let $\lambda \geq 0$. For a given Lipschitz continuous function $F : \mathcal{P}(K) \rightarrow \mathbb{R}$, we further define the functionals*

$$\begin{aligned} \mathcal{J}_n(G) &:= \begin{cases} \|G - F\|_{L^2(\mathcal{P}_2(K), \mathfrak{m}_n)}^2 + \lambda \mathfrak{pCE}_{2,\mathfrak{m}_n}(G), & \text{if } G \in \mathcal{V}_n, \\ +\infty & \text{if } G \in H^{1,2}(\mathcal{P}(K), W_2, \mathfrak{m}) \setminus \mathcal{V}_n, \end{cases} \\ \bar{\mathcal{J}}(G) &:= \|G - F\|_{L^2(\mathcal{P}_2(K), \mathfrak{m})}^2 + \lambda \mathfrak{CE}_{2,\mathfrak{m}}(G), \quad G \in H^{1,2}(\mathcal{P}(K), W_2, \mathfrak{m}). \end{aligned}$$

Then there exists a sequence of finite dimensional subsets of cylinder functions \mathcal{V}_n such that \mathcal{J}_n Γ -converges to $\bar{\mathcal{J}}$ as $n \rightarrow +\infty$. In particular, the minimizers of \mathcal{J}_n converge to the minimizer of $\bar{\mathcal{J}}$ as $n \rightarrow +\infty$ in $L^2(\mathcal{P}(K), \mathfrak{m})$. For $\lambda > 0$, the convergence is even in $H^{1,2}(\mathcal{P}(K), W_2, \mathfrak{m})$.

This result applies for instance, but not exclusively, to the Wasserstein distance function $F = F_\vartheta^{W_2}$. We further note that the above convergence is completely deterministic and it applies to any measure \mathfrak{m} ; i.e., it is not restricted to probability measures. However, it does not provide any convergence rate and it does not offer yet a proper analysis for data corrupted by noise. Moreover the subsets \mathcal{V}_n are not easily implementable as they require bounds on higher order derivatives, see Definition 4.1. For these reasons, we discuss in Section 4.2 a different approach, which leads to an analogous kind of convergence. The route of these results follows and generalizes [26]. However, it is based again on a randomization for the choice of \mathfrak{m}_n and we gain an explicit error estimate, albeit the convergence is only in mean and no longer deterministic. As a relevant element of novelty with respect to [26] that focused on pure least squares, let us also stress that we generalize the entire argument to Tikhonov-type functionals such as \mathcal{J}_n . Hence, we have to deal properly with the pre-Cheeger energy regularization term, which eventually contributes to control the propagation of the data noise on the final generalization error. (While we were concluding this paper, we came aware also of related results in the very recent preprint [88] that focuses again on a pure least squares rather than on a Tikhonov regularization model as we do here). Our main result in this direction is the following (see Theorem 4.12).

Theorem 1.4. *Let $K \subset \mathbb{R}^d$ be a compact set, \mathfrak{p} be a probability measure on $\mathcal{P}(K)$ and \mathfrak{p}_N , $N \in \mathbb{N}$, be the empirical measure associated to random points μ_1, \dots, μ_N which are i.i.d. with*

respect to \mathfrak{p} . Let $F : \mathcal{P}(K) \rightarrow [-M, M]$, $M > 0$, be a bounded and measurable function and let \tilde{F}^* be defined as

$$\tilde{F}^* := -M \vee S_N^{\lambda, n}(\tilde{F}) \wedge M,$$

where $S_N^{\lambda, n}(\tilde{F})$ is the minimizer of the functional

$$\mathcal{J}_{N, F}^\lambda(G) := \|\tilde{F} - G\|_{L^2(\mathcal{P}(K), \mathfrak{p}_N)}^2 + \lambda \mathfrak{pCE}_{2, \mathfrak{p}_N}(G), \quad G \in \mathcal{V}_n,$$

with \mathcal{V}_n , $n \in \mathbb{N}$, being an arbitrary n -dimensional subspace of cylinder functions, and

$$\tilde{F}(\mu_j) := F(\mu_j) + \eta_j, \quad j = 1, \dots, N,$$

is a noisy version of F obtained with i.i.d. random variables $(\eta_j)_{j=1}^N$ with variance bounded by $\sigma^2 > 0$. If $r > 0$ and M, n are suitably chosen (see (4.6)), then

$$\begin{aligned} \mathbb{E} \left[\|F - \tilde{F}^*\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 \right] &\leq 2e(F, n) \left(1 + \frac{c_{1/2}}{\log(N)(1+r)(1/2 + \lambda\mu_{\min, n})^2} \right) + 8\lambda \mathfrak{pCE}_{2, \mathfrak{p}_N}(P_n F) + \\ &+ 4 \frac{\sigma^2}{(1 + \lambda\mu_{\min, n})^2} \frac{n}{N} + 2M^2 N^{-r}, \end{aligned}$$

where $P_n F$ is the $L^2(\mathcal{P}(K), \mathfrak{p})$ -orthogonal projection of F onto \mathcal{V}_n ,

$$e(F, n) := \|P_n F - F\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2$$

is the \mathcal{V}_n -best approximation error with respect to the $L^2(\mathcal{P}(K), \mathfrak{p})$ -norm, and $\mu_{\min, n}$ is the minimal value of the pre-Cheeger energy evaluated on a suitable orthonormal basis of \mathcal{V}_n .

Up to scaling the parameter λ and choosing a subspace \mathcal{V}_n that entails a bound on the pre-Cheeger energy w.r.t. \mathfrak{p}_N , the above estimate ensures the convergence of \tilde{F}^* to F in expectation. Moreover, the explicit generalization error bound completely reveals the interplay between the number N of the noisy data, the dimension n of the space \mathcal{V}_n , the regularization parameter λ , and the noise level σ . In particular the contribution of the noise to the error bound can be controlled by either choosing a larger number N of data or by considering a larger regularization parameter λ that one needs to optimize. For $\lambda = 0$ one simply obtains again an analogous result as in [26].

As a final approach to the finite sample approximation problem, let us consider again the energy functional

$$\bar{\mathcal{J}}(G) := \left\| G - \tilde{F} \right\|_{L^2(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathfrak{m})}^2 + \lambda \mathfrak{CE}_{2, \mathfrak{m}}(G), \quad G \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathfrak{m}),$$

where \tilde{F} is a possibly noisy version of a given function $F \in L^2(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathfrak{m})$. Since $\bar{\mathcal{J}}$ is a strictly convex functional on $H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathfrak{m})$, its minimizer is equivalently the unique solution of the Euler–Lagrange equation specified in (5.3). Subsequently, stating the Euler–Lagrange equation as an operator equation in the dual space and employing the density of cylinder functions in the Wasserstein Sobolev space $H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathfrak{m})$ leads to a saddle point problem of the form

$$\inf_{G \in \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))} \sup_{\substack{H \in \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d)) \\ H \neq 0}} \frac{(G - \tilde{F}, H)_{L^2(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathfrak{m})} + \lambda \mathfrak{pCE}_{2, \mathfrak{m}}(G, H)}{|H|_{pH_{\mathfrak{m}}^{1,2}}},$$

where $\mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$ denotes again the space of all cylinder functions and $\mathfrak{pCE}_{2, \mathfrak{m}}(\cdot, \cdot)$ is defined as in (5.8) below. To solve this problem in applications, we have to replace the measure \mathfrak{m} by an empirical approximation \mathfrak{m}_N . Moreover, since we cannot efficiently numerically implement the entire set of all cylinder functions (especially in very high dimension), we restrict to a subset of cylinder functions represented by deep neural networks. In particular, similar as in [99], we

use an adversarial network for the computation of the supremum and a solution network for the infimum in the saddle point problem above; this gives rise to the adversarial training Algorithm 1.

We conclude our paper with a series of numerical experiments that test the three approaches illustrated above on the approximation of the Wasserstein distance between probability measures derived from datasets. Specifically, we employ two widely recognized datasets: MNIST, comprising handwritten characters, and CIFAR-10, featuring coloured 32×32 pixels images from 10 classes. In our approach, we interpret the elements within these datasets as probability measures. Consequently, we regard the normalized images within these datasets as representations of positive probability densities. First of all, we investigate how the approximation accuracy of (1.5) improves for an increasing number of precomputed potentials. Subsequently, we model the function (1.5) as a deep neural network, and examine how the error further decays if we optimize the network parameters by a suitable training procedure. Finally, we run some experiments for the solution of the Euler–Lagrange equation as briefly described above. We mention that the numerical experiments require considerable computational effort, which limits the number of training data that we consider and end with relatively large test errors. Yet, we obtain approximation errors, which are competitive in accuracy with respect to state of the art methods for computing the Wasserstein distance, with an improvement in computational time after training of several orders of magnitude, see Table 1.

Outline. In Section 2 we first introduce some basic notions and results about metric Sobolev spaces, Wasserstein spaces, and Wasserstein Sobolev spaces, which are fundamental for the present manuscript. Then, Section 3 deals with the explicit construction of cylinder functions that approximate the Wasserstein distance to a given reference measure, the random ε -subcovering of a complete and separable metric space, and the empirical approximation of the Wasserstein distance based on a finite number of samples of the distribution on the Wasserstein Sobolev space. Thereafter, we translate those results to the case of a discrete base space, which might be of interest in applications. In that setting, we also relate cylinder functions to deep neural networks. In Section 4.1, we verify that a sequence of cylinder functions obtained by empirical risk minimization converges to the sought solution, which could be the Wasserstein distance up to some arbitrarily small regularization term. In the subsequent Section 4.2, at the expense of determinism, we even obtain a convergence rate up to high probability of the empirical risk minimization approach. Then, in Section 5 we introduce a novel approach to compute the Wasserstein distance based on the Euler-Lagrange formulation of the risk minimization. The work is finally rounded off by some numerical experiments in the context of the MNIST and CIFAR-10 datasets.

Acknowledgments. M.F. and G.E.S. gratefully acknowledge the support of the Institute for Advanced Study of the Technical University of Munich, funded by the German Excellence Initiative. M.F. and P.H. acknowledge the support of the Munich Center for Machine Learning. The authors thank Felix Kraemer and Giuseppe Savaré for various useful conversations on the subject of this paper.

2. WASSERSTEIN SOBOLEV SPACES

2.1. The relaxation approach to Metric Sobolev spaces

In the last years, several notions of metric Sobolev spaces on metric measure spaces were studied, such as, e.g., the Newtonian approach in [84, 20]. Here we focus on the one presented by Ambrosio, Gigli and Savaré [4], where, starting from the ideas of Cheeger [24], they introduce the following notion of q -relaxed ($q \in (1, +\infty)$) gradient: given a metric measure space $(\mathcal{S}, \mathbf{d}, \mathbf{m})$ (this is to say that $(\mathcal{S}, \mathbf{d})$ is a complete and separable metric space and \mathbf{m} is a finite non-negative Borel measure on $(\mathcal{S}, \mathbf{d})$), a function $G \in L^q(\mathcal{S}, \mathbf{m})$ is a q -relaxed gradient (see [80, Definition 3.1.5])

of $f \in L^0(\mathcal{S}, \mathbf{m})$ if there exist a sequence of bounded Lipschitz functions $(f_n)_n \subset \text{Lip}_b(\mathcal{S}, \mathbf{d})$ and $\tilde{G} \in L^q(\mathcal{S}, \mathbf{m})$ such that

- (1) $f_n \rightarrow f$ in $L^0(\mathcal{S}, \mathbf{m})$ and $\text{lip}_d f_n \rightarrow \tilde{G}$ in $L^q(\mathcal{S}, \mathbf{m})$,
- (2) $\tilde{G} \leq G$ \mathbf{m} -a.e. in \mathcal{S} ;

here, for a function $h : \mathcal{S} \rightarrow \mathbb{R}$, the asymptotic Lipschitz constant is defined as

$$\text{lip}_d h(x) := \limsup_{y, z \rightarrow x, y \neq z} \frac{|h(y) - h(z)|}{\mathbf{d}(y, z)}, \quad x \in \mathcal{S}, \quad (2.1)$$

and $\text{Lip}_b(\mathcal{S}, \mathbf{d})$ denotes the space of \mathbf{d} -Lipschitz and bounded functions on \mathcal{S} . In particular, for $f \in \text{Lip}_b(\mathcal{S}, \mathbf{d})$, we have that

$$\text{Lip}(f, \mathcal{S}) := \sup_{x, y \in \mathcal{S}, x \neq y} \frac{|f(x) - f(y)|}{\mathbf{d}(x, y)} < \infty. \quad (2.2)$$

It is not difficult to see that if $f \in L^0(\mathcal{S}, \mathbf{m})$ admits at least one q -relaxed gradient, then it has a minimal one (both in L^q -norm and in the \mathbf{m} -a.e. sense) which is denoted by $|\mathbf{D}f|_{\star, q}$ and provides a notion of (norm of the) derivative for f . We can thus define the q -Cheeger energy of $f \in L^0(\mathcal{S}, \mathbf{m})$ as

$$\text{CE}_{q, \mathbf{m}}(f) := \begin{cases} \int_{\mathcal{S}} |\mathbf{D}f|_{\star, q}^q \, \mathbf{d}\mathbf{m} & \text{if } f \text{ admits a } q\text{-relaxed gradient,} \\ +\infty & \text{else.} \end{cases}$$

This quantity can be also obtained as the relaxation (see e.g. [80, Corollary 3.1.7]) of the so called pre- q -Cheeger energy

$$\text{pCE}_{q, \mathbf{m}}(f) := \int_{\mathcal{S}} (\text{lip}_d f)^q \, \mathbf{d}\mathbf{m}, \quad f \in \text{Lip}_b(\mathcal{S}, \mathbf{d}), \quad (2.3)$$

meaning that

$$\text{CE}_{q, \mathbf{m}}(f) = \inf \left\{ \liminf_{n \rightarrow +\infty} \text{pCE}_{q, \mathbf{m}}(f_n) : (f_n)_n \subset \text{Lip}_b(\mathcal{S}, \mathbf{d}), f_n \rightarrow f \text{ in } L^0(\mathcal{S}, \mathbf{m}) \right\}. \quad (2.4)$$

The resulting Sobolev space $H^{1, q}(\mathcal{S}, \mathbf{d}, \mathbf{m})$ is then the vector space of functions $f \in L^q(\mathcal{S}, \mathbf{m})$ with finite Cheeger energy endowed with the norm

$$\|f\|_{H^{1, q}(\mathcal{S}, \mathbf{d}, \mathbf{m})}^q := \int_{\mathcal{S}} |f|^q \, \mathbf{d}\mathbf{m} + \text{CE}_{q, \mathbf{m}}(f), \quad (2.5)$$

which makes it a Banach space (cf. [42, Theorem 2.1.17]). However, in general, even for $q = 2$, $H^{1, q}(\mathcal{S}, \mathbf{d}, \mathbf{m})$ is not a Hilbert space. For instance, for $\mathcal{S} = \mathbb{R}^d$, $\mathbf{d}(x, y) := |x - y|_\infty$ is the distance induced by the infinity norm, and \mathbf{m} is a Gaussian measure on \mathbb{R}^d , this is indeed not a Hilbert space.

2.2. Wasserstein spaces

We briefly collect here the main definitions related to the notion of Wasserstein spaces of probability measures. Given a metric space (U, ϱ) , we denote by $\mathcal{P}(U)$ the set of Borel probability measures on U and by $\mathbf{m}_{p, x_0}(\mu)$ the p -th moment, $p \in (1, +\infty)$, of a measure $\mu \in \mathcal{P}(U)$ defined as

$$\mathbf{m}_{p, x_0}(\mu) = \left(\int_U \varrho(x, x_0)^p \, \mathbf{d}\mu(x) \right)^{1/p}, \quad (2.6)$$

where $x_0 \in U$ is a fixed point. Usually U is a subset of \mathbb{R}^d and, in this case, we consider as distance ϱ the one induced by the Euclidean norm and we take as x_0 the origin in \mathbb{R}^d , also removing the subscript x_0 in the notation for the moment.

We consider the space $\mathcal{P}_p(U)$, which is the subset of $\mathcal{P}(U)$ of probability measures with finite p -th moment:

$$\mathcal{P}_p(U) := \{\mu \in \mathcal{P}(U) : m_{p,x_0}(\mu) < +\infty\}.$$

Notice that the above definition doesn't depend on the point x_0 that has been used. The p -Wasserstein distance between two points $\mu, \nu \in \mathcal{P}_p(U)$ is defined as

$$W_p(\mu, \nu) := \left(\inf \left\{ \int_{U \times U} \varrho(x, y)^p d\gamma(x, y) : \gamma \in \Gamma(\mu, \nu) \right\} \right)^{1/p},$$

where $\Gamma(\mu, \nu)$ denotes the subset of Borel probability measures on $U \times U$ having as marginals μ and ν , also called *transport plans* or simply *plans* between μ and ν .

It is well known that the infimum above is attained on a non empty and compact subset $\Gamma_o(\mu, \nu) \subset \Gamma(\mu, \nu)$. Notice that, in general, $\Gamma_o(\mu, \nu)$ also depends on p and ϱ , but we omit this dependence since it is always clear from the context which value of p and which distance ϱ we are referring to. Elements of $\Gamma_o(\mu, \nu)$ are called *optimal transport plans* or simply *optimal plans* between μ and ν and thus they satisfy

$$W_p^p(\mu, \nu) = \int_{U \times U} \varrho(x, y)^p d\gamma(x, y), \quad \text{for every } \gamma \in \Gamma_o(\mu, \nu). \quad (2.7)$$

It turns out that $(\mathcal{P}_p(U), W_p)$ is a metric space which is complete (resp. separable) if (U, ϱ) is complete (resp. separable), see [3, Proposition 7.1.5]. We also recall that $\mathcal{P}_p(U)$ is compact if and only if (U, ϱ) is compact (see [3, Remark 7.1.8]) and, in this case, we obviously have $\mathcal{P}_p(U) = \mathcal{P}(U)$ for every $p \in (1, +\infty)$.

We refer, e.g., to [96, 3, 95] for a comprehensive treatment of the theory of Optimal Transport and Wasserstein spaces.

2.3. Density of subalgebras of Lipschitz functions in Wasserstein Sobolev spaces

The starting point for the present work is the density result of [37]. Sticking for a short moment to the abstract framework of a general complete and separable metric space (\mathcal{S}, d) , it is shown in [37] that a sufficiently rich algebra of functions $\mathcal{A} \subset \text{Lip}_b(\mathcal{S}, d)$ is *dense in energy* in the Sobolev space $f \in H^{1,q}(\mathcal{S}, d, m)$. In particular, under suitable assumptions on the algebra \mathcal{A} , we have that for every $f \in H^{1,q}(\mathcal{S}, d, m)$ there exists a sequence $(f_n)_n \subset \mathcal{A}$ such that

$$f_n \rightarrow f \quad \text{and} \quad \text{lip}_d f_n \rightarrow |Df|_{*,q} \quad \text{in } L^q(\mathcal{S}, m). \quad (2.8)$$

Such a property is particularly relevant in case the algebra \mathcal{A} satisfies some interesting properties, which might not be clear for the whole set of bounded Lipschitz functions, and in turn can be transferred to the whole Sobolev space. For instance, this is the case for cylinder functions on the space of probability measures with the p -Wasserstein distance analyzed in [37] provided that the base space is a Hilbert space and $p = q = 2$, and subsequently further developed in [86] to the case of a Polish metric base space and $p, q \in (1, +\infty)$.

Let us recall the definition of cylinder function: a function $F : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ is a cylinder function if it is of the form

$$F = \psi \circ \mathbf{L}_\phi, \quad (2.9)$$

where $\psi \in C_b^1(\mathbb{R}^N)$, $\phi = (\phi_1, \phi_1, \dots, \phi_N) \in (C_b^1(\mathbb{R}^d))^N$, and $\mathbf{L}_\phi : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}^N$ being defined as

$$\mathbf{L}_\phi(\mu) := \left(\int_{\mathbb{R}^d} \phi_1 d\mu, \dots, \int_{\mathbb{R}^d} \phi_N d\mu \right), \quad \mu \in \mathcal{P}(\mathbb{R}^d). \quad (2.10)$$

The set of cylinder functions is denoted by $\mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$.

The ensuing result states that the asymptotic Lipschitz constant (cf. (2.1)) of a cylinder function has a simple expression, see [86, Proposition 4.7] and [37, Proposition 4.9] for the proof of the statement.

Proposition 2.1. *Let $F \in \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$ be as in (2.9) and let us define*

$$DF(\mu, x) := \sum_{n=1}^N \partial_n \psi(\mathbf{L}_\phi(\mu)) \nabla \phi_n(x), \quad (\mu, x) \in \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d. \quad (2.11)$$

Then, for every $\mu \in \mathcal{P}_p(\mathbb{R}^d)$, we have that

$$(\text{lip}_{W_p} F(\mu))^{p'} = \int_{\mathbb{R}^d} |DF(\mu, x)|^{p'} d\mu(x) =: \|DF[\mu]\|_{L_\mu^{p'}}^{p'}, \quad (2.12)$$

where $p' := p/(p-1)$.

The following statement is the density result mentioned above, see [37, Theorem 4.10] and [86, Theorem 4.15], and which is of relevance for the work presented herein.

Theorem 2.2. *Let \mathbf{m} be any non-negative and finite Borel measure on $(\mathcal{P}_p(\mathbb{R}^d), W_p)$. Then the algebra $\mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$ is dense in q -energy in $H^{1,q}(\mathcal{P}_p(\mathbb{R}^d), W_p, \mathbf{m})$, $p, q \in (1, +\infty)$. In particular, for every $F \in H^{1,q}(\mathcal{P}_p(\mathbb{R}^d), W_p, \mathbf{m})$ there exists a sequence $(F_n)_n \subset \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$ such that*

$$F_n \rightarrow F \text{ in } L^q(\mathcal{P}_p(\mathbb{R}^d), \mathbf{m}), \quad \rho \text{CE}_{q,\mathbf{m}}(F_n) \rightarrow \text{CE}_{q,\mathbf{m}}(F) \quad \text{as } n \rightarrow +\infty.$$

Moreover, the Banach space $H^{1,q}(\mathcal{P}_p(\mathbb{R}^d), W_p, \mathbf{m})$ is reflexive and uniformly convex for every $p, q \in (1, +\infty)$. Finally, if $p = q = 2$, then $\text{CE}_{2,\mathbf{m}}$ is a quadratic form and $H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m})$ is a Hilbert space.

Since we use it below, let us also emphasize that the structure of cylinder functions and the result above allow to extend the notion of vector-valued gradient to general functions in $H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m})$. In particular, we need below the properties summarized in the following remark.

Remark 2.3. For any $F, G \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m})$ there exist unique vector fields $D_{\mathbf{m}}F, D_{\mathbf{m}}G \in L^2(\mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d, \mathbf{m}; \mathbb{R}^d)$ such that

$$\text{CE}_{2,\mathbf{m}}(F, G) = \int_{\mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d} D_{\mathbf{m}}F \cdot D_{\mathbf{m}}G d\mathbf{m} = \int_{\mathcal{P}_2(\mathbb{R}^d)} \int_{\mathbb{R}^d} D_{\mathbf{m}}F(x, \mu) \cdot D_{\mathbf{m}}G(x, \mu) d\mu(x) d\mathbf{m}(\mu),$$

where $\mathbf{m} := \int_{\mathcal{P}_2(\mathbb{R}^d)} \delta_\mu d\mathbf{m}(\mu)$. For the proof see [37, Section 5]. Notice that even for $F \in \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$, in general, $D_{\mathbf{m}}F \neq DF$. This equality would correspond to the *closability* of the Dirichlet form induced by the Cheeger energy, which is not guaranteed in general. However it holds that

$$\int_{\mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d} D_{\mathbf{m}}F \cdot D_{\mathbf{m}}G d\mathbf{m} = \int_{\mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d} DF \cdot DG d\mathbf{m} \quad (2.13)$$

for every $F \in \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$ and $G \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m})$, see [37, Formula 5.36].

The above result, Theorem 2.2, has a twofold importance: from a theoretical point of view, in case $p = q = 2$, the Hilbertianity property is crucial, being the starting point for a very rich theory [70, 90, 91, 4, 40]. As a relevant example of application, in this work we make heavy use of the Hilbertianity of $H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m})$ in Section 5, where we compute the Euler-Lagrange equations of quadratic functionals over this space.

Moreover, again from an applied perspective, this theorem guarantees the approximability of a wide class of functions with the class of smooth cylinder functions. For instance, for any given $k \in (0, +\infty)$, it can be easily shown that the k -truncated p -Wasserstein distance from a fixed measure $\vartheta \in \mathcal{P}_p(\mathbb{R}^d)$,

$$\mu \mapsto W_p(\mu, \vartheta) \wedge k,$$

is an element of $H^{1,q}(\mathcal{P}_p(\mathbb{R}^d), W_p, \mathbf{m})$ for any non-negative Borel measure \mathbf{m} on $\mathcal{P}_p(\mathbb{R}^d)$; we thus know that there exists a sequence $(F_n)_n \subset \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$ such that

$$F_n \rightarrow W_p(\vartheta, \cdot) \wedge k \quad \text{and} \quad \text{lip}_{W_p} F_n \rightarrow |D(W_p(\vartheta, \cdot) \wedge k)|_{*,q} \quad \text{in } L^q(\mathcal{P}_p(\mathbb{R}^d), \mathbf{m}) \quad (2.14)$$

as $n \rightarrow +\infty$. Moreover, if the measure \mathbf{m} is concentrated on measures in $\mathcal{P}_p(\mathbb{R}^d)$ with the same compact support $K \subset \mathbb{R}^d$, (2.14) holds even with $k = +\infty$.

Let us remark that the result of Theorem 2.2 is not constructive, in the sense that in [37, 86] there is no explicit form for the sequence of cylinder functions approximating a given $F \in H^{1,q}(\mathcal{P}_p(\mathbb{R}^d), W_p, \mathbf{m})$. This is the content of the ensuing section.

3. APPROXIMABILITY OF WASSERSTEIN DISTANCES BY CYLINDER MAPS

3.1. Distribution independent approximation

We first present a general strategy for the pointwise approximation of the p -Wasserstein distance through Kantorovich potentials, which is similar to the one proposed in [31, 37]. In particular we show that the Wasserstein distance $W_p(\cdot, \vartheta)$, for a given reference measure $\vartheta \in \mathcal{P}_p(K) = \mathcal{P}(K)$ with $K \subset \mathbb{R}^d$ being compact, can be approximated pointwise by a well-constructed uniformly bounded sequence of cylinder functions $(F_n)_n \subset \mathfrak{C}(\mathcal{P}(K), C_b^1(\mathbb{R}^d))$; i.e., we have that

$$\lim_{n \rightarrow \infty} F_n(\mu) = W_p(\mu, \vartheta). \quad (3.1)$$

As a consequence of the pointwise convergence and dominated convergence theorem, for any non-negative Borel measure \mathbf{m} on $\mathcal{P}(K)$ and any $q \in [1, +\infty)$, we obtain the convergence in $L^q(\mathcal{P}(K), \mathbf{m})$:

$$\lim_{n \rightarrow \infty} \int_{\mathcal{P}(K)} |F_n(\mu) - W_p(\mu, \vartheta)|^q d\mathbf{m}(\mu) = 0. \quad (3.2)$$

We start reporting a result (the proof of which can be found in [86]) which is the fundamental link between Wasserstein distance and cylinder functions. In the following, we denote by $\mathcal{P}_p^r(\mathbb{R}^d)$ the subset of $\mathcal{P}_p(\mathbb{R}^d)$ of *regular measures*; i.e., those elements $\mu \in \mathcal{P}_p(\mathbb{R}^d)$ such that $\mu \ll \mathcal{L}^d$, where \mathcal{L}^d denotes the Lebesgue measure on \mathbb{R}^d .

Theorem 3.1. *Let $\mu, \nu \in \mathcal{P}_p^r(\mathbb{R}^d)$ with $\text{supp } \nu = \overline{B(0, R)}$ for some $R > 0$, where $B(0, R)$ denotes the Euclidean ball of radius R centered at the origin. There exists a unique pair of locally Lipschitz functions $\phi \in L^1(B(0, R), \nu)$ and $\psi \in L^1(\mathbb{R}^d, \mu)$ such that*

$$(i) \quad \phi(x) + \psi(y) \leq \frac{1}{p}|x - y|^p \quad \text{for every } (x, y) \in B(0, R) \times \mathbb{R}^d,$$

(ii)

$$\phi(x) = \inf_{y \in \mathbb{R}^d} \left\{ \frac{1}{p}|x - y|^p - \psi(y) \right\} \quad \text{for every } x \in B(0, R), \quad (3.3)$$

$$\psi(y) = \inf_{x \in B(0, R)} \left\{ \frac{1}{p}|x - y|^p - \phi(x) \right\} \quad \text{for every } y \in \mathbb{R}^d, \quad (3.4)$$

$$(iii) \quad \psi(0) = 0,$$

$$(iv) \quad \int_{B(0, R)} \phi d\nu + \int_{\mathbb{R}^d} \psi d\mu = \frac{1}{p} W_p^p(\mu, \nu).$$

Such a unique pair is denoted by $(\Phi(\nu, \mu), \Phi^*(\nu, \mu))$. Finally, the function $\psi := \Phi^*(\nu, \mu)$ satisfies the following estimates: there exists a constant $K_{p,R}$, depending only on p and R , such that

$$|\psi(y') - \psi(y'')| \leq |y' - y''| 2^{p-1} (2R^{p-1} + |y'|^{p-1} + |y''|^{p-1}) \quad \text{for every } y', y'' \in \mathbb{R}^d, \quad (3.5)$$

$$|\psi(y)| \leq K_{p,R} (1 + |y|^p) \quad \text{for every } y \in \mathbb{R}^d. \quad (3.6)$$

Remark 3.2. There exists a constant $D_{p,R}$ depending only on p and R such that, if $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$ is a function satisfying (3.5), then

$$\int_{\mathbb{R}^d} \psi \, d(\mu' - \mu'') \leq D_{p,R} W_p(\mu', \mu'') (1 + \mathfrak{m}_p(\mu') + \mathfrak{m}_p(\mu'')) \quad \text{for every } \mu', \mu'' \in \mathcal{P}_p(\mathbb{R}^d), \quad (3.7)$$

where $\mathfrak{m}_p(\cdot)$ is as in (2.6). Indeed, by (3.5), if we take any optimal plan $\gamma \in \Gamma_o(\mu', \mu'')$ as in (2.7) (for the distance induced by the Euclidean norm), we have

$$\begin{aligned} \int_{\mathbb{R}^d} \psi \, d(\mu' - \mu'') &= \int_{\mathbb{R}^d \times \mathbb{R}^d} (\psi(y') - \psi(y'')) \, d\gamma(y', y'') \\ &\leq 2^{p-1} \left(\int_{\mathbb{R}^d \times \mathbb{R}^d} |y' - y''|^p \, d\gamma(y', y'') \right)^{1/p} \\ &\quad \left(\int_{\mathbb{R}^d \times \mathbb{R}^d} (2R^{p-1} + |y'|^{p-1} + |y''|^{p-1})^{p'} \, d\gamma(y', y'') \right)^{1/p'} \end{aligned}$$

so that (3.7) follows.

Remark 3.3. Let $\{\psi_h\}_{h=1}^n$, $n \in \mathbb{N}_0$, be a set of functions satisfying (3.5) and let

$$G_n(\mu) := \max_{1 \leq h \leq n} \int_{\mathbb{R}^d} \psi_h \, d\mu, \quad \mu \in \mathcal{P}_p(\mathbb{R}^d).$$

Then

$$|G_n(\mu') - G_n(\mu'')| \leq D_{p,R} W_p(\mu', \mu'') (1 + \mathfrak{m}_p(\mu') + \mathfrak{m}_p(\mu'')) \quad \text{for every } \mu', \mu'' \in \mathcal{P}_p(\mathbb{R}^d),$$

where $D_{p,R}$ is the constant from Remark 3.2. Indeed, by Remark 3.2 and the definition of G_n we have that

$$\int_{\mathbb{R}^d} \psi_h \, d\mu' - G_n(\mu'') \leq \int_{\mathbb{R}^d} \psi_h \, d\mu' - \int_{\mathbb{R}^d} \psi_h \, d\mu'' \leq D_{p,R} W_p(\mu', \mu'') (1 + \mathfrak{m}_p(\mu') + \mathfrak{m}_p(\mu''))$$

for every $1 \leq h \leq n$. Hence, passing to the maximum in h we obtain that

$$G_n(\mu') - G_n(\mu'') \leq D_{p,R} W_p(\mu', \mu'') (1 + \mathfrak{m}_p(\mu') + \mathfrak{m}_p(\mu'')).$$

Reversing the roles of μ' and μ'' gives the sought inequality.

The main result of this subsection is a slightly more explicit and constructive version of the approximation results contained in [31, 37]. Before stating it, we need to fix some notation.

Let $\kappa \in C_c^\infty(\mathbb{R}^d)$ be such that $\text{supp } \kappa = \overline{B(0, 1)}$, $\kappa(x) \geq 0$ for every $x \in \mathbb{R}^d$, $\kappa(x) > 0$ for every $x \in B(0, 1)$, $\int_{\mathbb{R}^d} \kappa \, d\mathcal{L}^d = 1$, and $\kappa(-x) = \kappa(x)$ for every $x \in \mathbb{R}^d$. Then, for any $0 < \varepsilon < 1$, we consider the standard mollifiers

$$\kappa_\varepsilon(x) := \frac{1}{\varepsilon^d} \kappa(x/\varepsilon), \quad x \in \mathbb{R}^d.$$

Given $\sigma \in \mathcal{P}_p(\mathbb{R}^d)$, $0 < \varepsilon < 1$, and $R > 0$, we further define

$$\sigma_\varepsilon := \sigma * \kappa_\varepsilon, \quad \hat{\sigma}_{\varepsilon,R} := \frac{\sigma_\varepsilon \llcorner B(0, R) + \varepsilon^{d+p+1} \mathcal{L}^d \llcorner B(0, R)}{\sigma_\varepsilon(B(0, R)) + \varepsilon^{d+p+1} \mathcal{L}^d(B(0, R))};$$

here, $*$ and \lfloor are the convolution and restriction operators, respectively. Notice that $\sigma_\varepsilon, \hat{\sigma}_{\varepsilon,R} \in \mathcal{P}_p^r(\mathbb{R}^d)$ with

$$\text{supp } \sigma_\varepsilon \subset \text{supp } \sigma + \text{B}(0, \varepsilon), \quad \text{supp } \hat{\sigma}_{\varepsilon,R} = \overline{\text{B}(0, R)}, \quad \hat{\sigma}_\varepsilon \geq \frac{\varepsilon^{d+p+1}}{1 + \varepsilon^{d+p+1}\omega_{R,d}} \mathcal{L}^d \lfloor \text{B}(0, R),$$

where $\omega_{R,d} := \mathcal{L}^d(\text{B}(0, R))$. We have also that $W_p(\sigma, \sigma_\varepsilon) \leq \varepsilon \mathfrak{m}_p(\kappa \mathcal{L}^d)$ (see [3, Lemma 7.1.10]) so that $W_p(\sigma_\varepsilon, \sigma) \rightarrow 0$ as $\varepsilon \rightarrow 0$. Moreover, if $\sigma, \sigma' \in \mathcal{P}_p(\mathbb{R}^d)$, we have (cf. [79, Lemma 5.2])

$$W_p(\sigma_\varepsilon, \sigma'_\varepsilon) \leq W_p(\sigma, \sigma') \quad \text{for every } 0 < \varepsilon < 1. \quad (3.8)$$

We can also see that $W_p(\hat{\sigma}_{\varepsilon,1/\varepsilon}, \sigma) \rightarrow 0$ as $\varepsilon \rightarrow 0$, which follows from the well known fact (cf. [3, Proposition 7.1.5]) that the convergence in W_p is equivalent to the convergence

$$\int_{\mathbb{R}^d} \varphi \, d\hat{\sigma}_{\varepsilon,R} \rightarrow \int_{\mathbb{R}^d} \varphi \, d\sigma$$

for every continuous function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ with less than p -growth. This condition is easily seen to be satisfied in our case, also noticing that $\sigma_\varepsilon(\text{B}(0, 1/\varepsilon)) \rightarrow 1$: indeed for every $\delta > 0$ we can find $R_\delta > 0$ such that $\sigma(\text{B}(0, R_\delta)) > 1 - \delta$ so that, if $\varepsilon < 1/R_\delta$, we have

$$\sigma_\varepsilon(\text{B}(0, 1/\varepsilon)) \geq \sigma_\varepsilon(\text{B}(0, R_\delta))$$

hence we get

$$\liminf_{\varepsilon \downarrow 0} \sigma_\varepsilon(\text{B}(0, 1/\varepsilon)) \geq \liminf_{\varepsilon \downarrow 0} \sigma_\varepsilon(\text{B}(0, R_\delta)) \geq \sigma(\text{B}(0, R_\delta)) > 1 - \delta$$

and being $\delta > 0$ arbitrary we conclude.

In case σ has compact support, we can also obtain an explicit rate of convergence of $\hat{\sigma}_{\varepsilon,R}$ to σ for a particular choice of R : we can take any $R > 1$ such that $\text{supp } \sigma \subset \text{B}(0, R-1)$ obtaining thus that $\text{supp } \sigma_\varepsilon \subset \text{B}(0, R)$ and hence that $\sigma_\varepsilon \lfloor \text{B}(0, R) = \sigma_\varepsilon$. Using the convexity of the Wasserstein distance and the triangle inequality, it is not difficult to see that

$$W_p(\hat{\sigma}_{\varepsilon,R}, \sigma) \leq \varepsilon \mathfrak{m}_p(\kappa \mathcal{L}^d) + \varepsilon^{d+p+1} \omega_{R,d} (\mathfrak{m}_p(\sigma) + R) \leq (1 + 2\omega_{R,d} R) \varepsilon. \quad (3.9)$$

Proposition 3.4. *Let $K \subset \mathbb{R}^d$ be a compact set, $\vartheta \in \mathcal{P}_p(K) = \mathcal{P}(K)$, and $(\sigma^h)_h \subset \mathcal{P}(K)$ be a dense subset in $\mathcal{P}(K)$ with respect to the p -Wasserstein distance. Moreover, let $R > 1$ be such that $K \subset \text{B}(0, R-1)$, and for every $n, h \in \mathbb{N}$ let us consider the functions*

$$\varphi_n^h := \Phi(\hat{\vartheta}_{1/n,R}, \sigma_{1/n}^h), \quad \psi_n^h := \Phi^*(\hat{\vartheta}_{1/n,R}, \sigma_{1/n}^h), \quad u_n^h(x) := \psi_n^h(x) + \int_{\mathbb{R}^d} \varphi_n^h \, d\hat{\vartheta}_{1/n,R}, \quad x \in \mathbb{R}^d,$$

where Φ and Φ^* are as in Theorem 3.1. We further define, for every $k \in \mathbb{N}$, the compact sets

$$C_n := [-\|u_n^1 * \kappa_{1/n}\|_\infty, \|u_n^1 * \kappa_{1/n}\|_\infty] \times \cdots \times [-\|u_n^n * \kappa_{1/n}\|_\infty, \|u_n^n * \kappa_{1/n}\|_\infty] \subset \mathbb{R}^k,$$

where $\|\cdot\|_\infty$ denotes the infinity norm on K , and choose a smooth increasing approximation $\eta_n : \mathbb{R}^n \rightarrow \mathbb{R}$ of the function

$$(x_1, \dots, x_n) \mapsto \left(p \max_{1 \leq h \leq n} x_h \right)^p$$

such that

$$\sup_{(x_1, \dots, x_n) \in C_n} \left| \eta_n(x_1, \dots, x_n) - \left(p \max_{1 \leq h \leq n} x_h \right)^{1/p} \right| \leq 1/n. \quad (3.10)$$

Then, the cylinder map

$$F_n(\vartheta, \mu) := \eta_n \left(\int_{\mathbb{R}^d} (u_n^1 * \kappa_{1/n}) \, d\mu, \int_{\mathbb{R}^d} (u_n^2 * \kappa_{1/n}) \, d\mu, \dots, \int_{\mathbb{R}^d} (u_n^n * \kappa_{1/n}) \, d\mu \right), \quad \mu \in \mathcal{P}(K),$$

converges pointwise from below to $W_p(\vartheta, \cdot)$ in $\mathcal{P}(K)$.

Proof. Let us define

$$G_n(\mu) := \max_{1 \leq h \leq n} \int_{\mathbb{R}^d} u_n^h d\mu_{1/n}, \quad \mu \in \mathcal{P}(K)$$

and observe that, thanks to Theorem 3.1,

$$G_n(\sigma^h) = \frac{1}{p} W_p^p(\hat{\vartheta}_{1/n,R}, \sigma_{1/n}^h) \quad \text{for every } 1 \leq h \leq n.$$

In contrast, for the cylinder function F_n we only have that $F_n \leq W_p(\cdot, \vartheta)$. Let us fix an arbitrary $\mu \in \mathcal{P}(K)$ and $n \in \mathbb{N}$; by employing the triangle inequality, we find that

$$\begin{aligned} |F_n(\vartheta, \mu) - W_p(\mu, \vartheta)| &\leq \left| F_n(\vartheta, \mu) - \left(p \max_{1 \leq h \leq n} \int_{\mathbb{R}^d} (u_n^h * \kappa_{1/n}) d\mu \right)^{1/p} \right| \\ &\quad + \left| (pG_n(\mu))^{1/p} - (pG_n(\sigma^h))^{1/p} \right| + \left| W_p(\hat{\vartheta}_{1/n,R}, \sigma_{1/n}^h) - W_p(\hat{\vartheta}_{1/n,R}, \mu_{1/n}) \right| \\ &\quad + \left| W_p(\hat{\vartheta}_{1/n,R}, \mu_{1/n}) - W_p(\hat{\vartheta}_{1/n,R}, \mu) \right| + \left| W_p(\hat{\vartheta}_{1/n,R}, \mu) - W_p(\vartheta, \mu) \right|. \end{aligned}$$

Invoking (3.10) and (3.8), together with the triangle inequality, further leads to

$$\begin{aligned} |F_n(\vartheta, \mu) - W_p(\mu, \vartheta)| &\leq 1/n + p^{1/p} \left| G_n(\mu) - G_n(\sigma^h) \right|^{1/p} + W_p(\sigma^h, \mu) \\ &\quad + W_p(\mu_{1/n}, \mu) + W_p(\hat{\vartheta}_{1/n,R}, \vartheta). \end{aligned}$$

Hence, by Remark 3.3, we arrive at

$$\begin{aligned} |F_n(\vartheta, \mu) - W_p(\mu, \vartheta)| &\leq 1/n + p^{1/p} D_{p,R}^{1/p} W_p^{1/p}(\mu, \sigma^h) (1 + \mathfrak{m}_p(\mu) + \mathfrak{m}_p(\sigma^h))^{1/p} \\ &\quad + W_p(\sigma^h, \mu) + W_p(\mu_{1/n}, \mu) + W_p(\hat{\vartheta}_{1/n,R}, \vartheta). \end{aligned}$$

Passing first to the limit as $n \rightarrow \infty$ and then taking the infimum with respect to $h \in \mathbb{N}$ implies the claim. \square

On the one hand, the pointwise convergence (3.1) is a deterministic and universal result, because it renders (3.2) valid for any suitable \mathfrak{m} and it does not depend on \mathfrak{m} . On the other hand, it does not provide any rate of convergence. In the following we shall trade deterministic and universal approximation with more control on the rate of convergence. For that we need to introduce the concept of random subcovering.

3.2. Random subcoverings

In this section we introduce random subcoverings in the general context of a complete and separable metric space (\mathcal{S}, d) endowed with a Borel probability measure \mathfrak{p} (we use the notation \mathfrak{p} to distinguish it from a general non-negative Borel measure \mathfrak{m}). We say in short that $(\mathcal{S}, d, \mathfrak{p})$ is a Polish metric-probability space. We often deal with points X_1, \dots, X_n , $n \in \mathbb{N}$, in \mathcal{S} drawn randomly according to \mathfrak{p} , meaning that we have fixed a probability space $(\Omega, \mathcal{E}, \mathbb{P})$ and $(X_i)_\# \mathbb{P} = \mathfrak{p}$ for $i = 1, \dots, n$, where $\#$ denotes the push forward operator. We recall that, for measurable spaces (E_1, \mathcal{E}_1) and (E_2, \mathcal{E}_2) , if μ is a measure on (E_1, \mathcal{E}_1) and $f : E_1 \rightarrow E_2$ is $\mathcal{E}_1 - \mathcal{E}_2$ measurable, the measure $f_\# \mu$ on (E_2, \mathcal{E}_2) is defined as

$$(f_\# \mu)(B) := \mu(f^{-1}(B)), \quad B \in \mathcal{E}_2.$$

We also use the notation \mathbb{E} to denote the expected value on $(\Omega, \mathcal{E}, \mathbb{P})$; i.e., the integral w.r.t. the probability measure \mathbb{P} .

Definition 3.5. We introduce the following definitions.

- For $\varepsilon > 0$ and $n \in \mathbb{N}$, we define a *random ε -subcovering of size n of $(\mathcal{S}, d, \mathfrak{p})$* as a finite collection of n radius ε -balls $\mathcal{C} = \{B(X_i, \varepsilon) : i = 1, \dots, n\}$, whose centers X_1, \dots, X_n are drawn i.i.d. at random according to \mathfrak{p} .

- For $\varepsilon > 0$ and $n \in \mathbb{N}$, we call (ε, n) -subcovering probability the probability

$$p_{\varepsilon, n} := \mathbb{P}\left(X \in \cup_{i=1}^n B(X_i, \varepsilon)\right),$$

where X, X_1, \dots, X_n are drawn i.i.d. according to \mathbf{p} .

- Finally, for $\delta, \varepsilon > 0$ we define the ε -subcovering number $N_{\delta, \varepsilon}(\mathcal{S}, \mathbf{d}, \mathbf{p})$ as the minimal size of random (ε, n) -subcoverings to cover \mathcal{S} with probability $1 - \delta$, i.e.,

$$N_{\delta, \varepsilon}(\mathcal{S}, \mathbf{d}, \mathbf{p}) := \min \{n \in \mathbb{N} : p_{\varepsilon, n} \geq 1 - \delta\}.$$

Notice that, while for a non-compact metric space \mathcal{S} the classical deterministic covering number is not finite, the (δ, ε) -subcovering number may be instead finite and small. For instance, if the support of the measure \mathbf{p} is itself contained in a ball $B(\bar{x}, \varepsilon/2)$, then obviously $p_{\varepsilon, n} = 1$ for all $n \in \mathbb{N}$, and therefore $N_{\delta, \varepsilon}(\mathcal{S}) = 1$. The ε -subcovering probabilities and relative (δ, ε) -subcovering number essentially measure how locally concentrated the measure \mathbf{p} is. To some extent it is a form of quantification of the entropy of the measure. Let us now compute explicitly the ε -subcovering probabilities $p_{\varepsilon, n}$ and estimate the (δ, ε) -subcovering number, respectively.

Proposition 3.6. *Let $(\mathcal{S}, \mathbf{d}, \mathbf{p})$ be a Polish metric-probability space. If $n \in \mathbb{N}$ and $\varepsilon > 0$, then we have*

$$p_{\varepsilon, n} = 1 - \int_{\mathcal{S}} (1 - \mathbf{p}(B(x, \varepsilon)))^n d\mathbf{p}(x) = 1 - \mathbb{E}[\mathbf{p}(\cap_{i=1}^n B^c(X_i, \varepsilon))]. \quad (3.11)$$

As a consequence the (δ, ε) -subcovering number can be estimated from above as

$$N_{\delta, \varepsilon}(\mathcal{S}, \mathbf{d}, \mathbf{p}) \leq \left\lceil \frac{\log(\delta)}{\int_{\mathcal{S}} \log(\mathbf{p}(B^c(x, \varepsilon))) d\mathbf{p}(x)} \right\rceil. \quad (3.12)$$

Proof. Let us first consider

$$p_{\varepsilon, n} = 1 - \mathbb{P}\left(X \in \bigcap_{i=1}^n B^c(X_i, \varepsilon)\right),$$

and compute

$$\begin{aligned} \mathbb{P}\left(X \in \bigcap_{i=1}^n B^c(X_i, \varepsilon)\right) &= \int_{\mathcal{S}^{n+1}} \chi_{\{|x-x_i| \geq \varepsilon \forall i=1, \dots, n\}}(x, x_1, \dots, x_n) d(\otimes^{n+1} \mathbf{p})(x, x_1, \dots, x_n) \\ &= \int_{\mathcal{S}} \int_{\mathcal{S}^n} \chi_{\{|x-x_i| \geq \varepsilon \forall i=1, \dots, n\}}(x, x_1, \dots, x_n) d(\otimes^n \mathbf{p})(x_1, \dots, x_n) d\mathbf{p}(x) \\ &= \int_{\mathcal{S}} \prod_{i=1}^n \mathbb{P}(X_i \in B^c(x, \varepsilon)) d\mathbf{p}(x) \\ &= \int_{\mathcal{S}} (1 - \mathbf{p}(B(x, \varepsilon)))^n d\mathbf{p}(x), \end{aligned}$$

where we have denoted by $\otimes^m \mathbf{p}$ the product measure on \mathcal{S}^m obtained multiplying m copies of \mathbf{p} , for $m = n, n + 1$. Thus (3.11) is verified. In turn, this further shows that $p_{\varepsilon, n} \geq 1 - \delta$ if and only if

$$\int_{\mathcal{S}} (1 - \mathbf{p}(B(x, \varepsilon)))^n d\mathbf{p}(x) = \mathbb{P}\left(X \in \bigcap_{i=1}^n B^c(X_i, \varepsilon)\right) \leq \delta.$$

By applying the logarithm on both sides and using Jensen's inequality we obtain

$$n \geq \frac{\log(\delta)}{\int_{\mathcal{S}} \log(\mathbf{p}(B^c(x, \varepsilon))) d\mathbf{p}(x)}.$$

Hence, one deduces (3.12). \square

Lemma 3.7. *Let $(\mathcal{S}, d, \mathbf{p})$ be a Polish metric-probability space. We have that $0 \leq p_{\varepsilon, n} \leq 1$, $p_{\varepsilon, 0} = 0$, $p_{\varepsilon, 1} = \int_{\mathcal{S}} \mathbf{p}(B(x, \varepsilon)) d\mathbf{p}(x)$, and $\lim_{n \rightarrow \infty} p_{\varepsilon, n} = 1$. Moreover, $\text{supp } \mathbf{p}$ is compact if and only if*

$$\underline{p}_{\varepsilon} := \inf_{x \in \text{supp } \mathbf{p}} \mathbf{p}(B(x, \varepsilon)) > 0 \quad \text{for every } \varepsilon > 0. \quad (3.13)$$

In this case, we have that

$$p_{\varepsilon, n} \geq 1 - (1 - \underline{p}_{\varepsilon})^n, \quad (3.14)$$

and $p_{\varepsilon, n}$ converges exponentially fast to 1 for $n \rightarrow \infty$.

Proof. If $x \in \text{supp } \mathbf{p}$, then, for any $\varepsilon > 0$, we have that $\mathbf{p}(B(x, \varepsilon)) > 0$ and in turn

$$\lim_{n \rightarrow \infty} (1 - \mathbf{p}(B(x, \varepsilon)))^n = 0.$$

Consequently, by (3.11) and the dominated convergence theorem we deduce that $\lim_{n \rightarrow \infty} p_{\varepsilon, n} = 1$. Assume now that $\text{supp } \mathbf{p}$ is compact, then (3.13) is satisfied because of the lower semicontinuity of the map $x \mapsto \mathbf{p}(B(x, \varepsilon))$. On the other hand, we prove that if $\text{supp } \mathbf{p}$ is not compact, then (3.13) is always violated. So let us assume that $\text{supp } \mathbf{p}$ is not compact and $\underline{p}_{\varepsilon} > 0$. Since $\text{supp } \mathbf{p}$ is not compact there exists some $\varepsilon > 0$ such that any finite union of balls with radii equal to ε does not cover $\text{supp } \mathbf{p}$. Moreover, since \mathbf{p} is tight, we can find a compact $K \subset \text{supp } \mathbf{p}$ such that $\mathbf{p}(\text{supp } \mathbf{p} \setminus K) < \underline{p}_{\varepsilon/2}$. In turn, by the compactness of K , there exists an integer $N \in \mathbb{N}$ and points $x_1, \dots, x_N \in K$ such that

$$K \subset \bigcup_{i=1}^N B(x_i, \varepsilon/2).$$

Let $x \in \text{supp } \mathbf{p} \setminus \bigcup_{i=1}^N B(x_i, \varepsilon)$. Then $d(x, x_i) \geq \varepsilon$ for every $i = 1, \dots, N$ and thus $B(x, \varepsilon/2) \cap K = \emptyset$. But this means that $\mathbf{p}(B(x, \varepsilon/2)) < \underline{p}_{\varepsilon/2}$, a contradiction. \square

Remark 3.8. We observe that $p_{\varepsilon, 1} = \int_{\mathcal{S}} \mathbf{p}(B(x, \varepsilon)) d\mathbf{p}(x)$ is a form of regularized negative entropy of the probability measure \mathbf{p} , where the density of \mathbf{p} is estimated by $\mathbf{p}(B(x, \varepsilon))$. Further notice that it may be that $\underline{p}_{\varepsilon} \rightarrow 0$ very fast for $\varepsilon \rightarrow 0$. Hence, the rate of convergence (3.14) has a meaning as long as $\varepsilon > 0$ is fixed and not too small. In some applications, as the ones we consider below, one may not need to reach arbitrary small accuracy, but be content with a sufficiently small approximation.

The following result shows that (ε, n) -subcovering probabilities depend continuously on the distribution.

Lemma 3.9. *Let $(\mathcal{S}, d, \mathbf{p})$ be a Polish metric-probability space. Assume that $\mathbf{p}^k \ll \mathbf{p}$, i.e., \mathbf{p}^k is absolutely continuous with respect to \mathbf{p} , and $\rho_k := \frac{d\mathbf{p}^k}{d\mathbf{p}} \rightarrow 1$ in $L^1(\mathcal{S}, \mathbf{p})$ as $k \rightarrow \infty$. Then*

$$\lim_{k \rightarrow \infty} \int_{\mathcal{S}} \varphi(x) d\mathbf{p}^k(x) = \int_{\mathcal{S}} \varphi(x) d\mathbf{p}(x), \quad (3.15)$$

for any $\varphi \in L^\infty(\mathcal{S})$. Hence, \mathbf{p}^k converges narrowly to \mathbf{p} . Additionally, for any fixed $\varepsilon > 0$ and $n \in \mathbb{N}$ we have

$$\lim_{k \rightarrow \infty} p_{\varepsilon, n}(\mathbf{p}^k) = p_{\varepsilon, n}(\mathbf{p}). \quad (3.16)$$

Proof. The first limit follows from

$$\left| \int_{\mathcal{S}} \varphi(x) (\rho_k(x) - 1) d\mathbf{p}(x) \right| \leq \|\varphi\|_\infty \int_{\mathcal{S}} |\rho_k(x) - 1| d\mathbf{p}(x) \rightarrow 0, \quad k \rightarrow \infty.$$

Concerning (3.16), employing (3.11) and the triangle inequality yields that

$$\begin{aligned}
 |p_{\varepsilon,n}(\mathbf{p}^k) - p_{\varepsilon,n}(\mathbf{p})| &\leq \int_{\mathcal{S}} \left(\int_{\mathcal{S}} \chi_{B(x,\varepsilon)^c}(y) \rho_k(y) d\mathbf{p}(y) \right)^n |\rho_k(x) - 1| d\mathbf{p}(x) \\
 &\quad + \int_{\mathcal{S}} \left| \left(\int_{\mathcal{S}} \chi_{B(x,\varepsilon)^c}(y) \rho_k(y) d\mathbf{p}(y) \right)^n - \mathbf{p}(B(x,\varepsilon)^c)^n \right| d\mathbf{p}(x) \\
 &\leq \|\rho_k - 1\|_{L^1(\mathcal{S},\mathbf{p})} \\
 &\quad + \int_{\mathcal{S}} \left| \left(\int_{\mathcal{S}} \chi_{B(x,\varepsilon)^c}(y) \rho_k(y) d\mathbf{p}(y) \right)^n - \mathbf{p}(B(x,\varepsilon)^c)^n \right| d\mathbf{p}(x).
 \end{aligned}$$

The first term converges to 0 as $k \rightarrow \infty$ by assumption. Moreover, the second term vanishes as well, as $k \rightarrow \infty$, by the pointwise convergence ensured by (3.15) and the dominated convergence theorem. \square

We present now some applications of (ε, n) -subcovering probabilities in estimating integrals by means of empirical measures. In the results below we presume that it is possible to estimate the measure \mathbf{p} on suitable sets C_i for the sake of allocating suitable weights of quadrature formulas. Below we consider the space of bounded Lipschitz continuous functions $\text{Lip}_b(\mathcal{S}, \mathbf{d})$ with norm $\|\varphi\|_{BL} = \|\varphi\|_{\infty} + \text{Lip}(\varphi, \mathcal{S})$, where $\text{Lip}(\varphi, \mathcal{S})$ is the Lipschitz constant of φ (cf. (2.2)), and its dual distance

$$d_{BL}(\mathbf{p}, \mathbf{p}') = \sup_{\varphi \in \text{Lip}_b(\mathcal{S}, \mathbf{d}), \|\varphi\|_{BL} \leq 1} \left| \int_{\mathcal{S}} \varphi(x) d\mathbf{p}(x) - \int_{\mathcal{S}} \varphi(x) d\mathbf{p}'(x) \right|.$$

Lemma 3.10. *Let $(\mathcal{S}, \mathbf{d}, \mathbf{p})$ be a Polish metric-probability space. Given a random (ε, n) -subcovering $\mathcal{C} = \{B(X_i, \varepsilon) : i = 1, \dots, n\}$ of \mathcal{S} , we consider disjoint sets $C_i \subset B(X_i, \varepsilon)$, $i = 1, \dots, n$, such that $\cup_{i=1}^n C_i = \cup_{i=1}^n B(X_i, \varepsilon)$ and we fix the empirical measure*

$$\mathbf{p}^{n,\varepsilon} = \frac{1}{\sum_{i=1}^n \mathbf{p}(C_i)} \sum_{i=1}^n \mathbf{p}(C_i) \delta_{X_i}. \quad (3.17)$$

Then, we have that

$$\mathbb{E} [d_{BL}(\mathbf{p}, \mathbf{p}^{n,\varepsilon})] \leq \varepsilon p_{\varepsilon,n} + 2(1 - p_{\varepsilon,n}). \quad (3.18)$$

Proof. For any $x \in \cup_{i=1}^n B(X_i, \varepsilon)$ we associate the unique $i = i(x)$ such that $x \in C_i$ and define $X^*(x) := X_{i(x)}$. With this notation we may write

$$\begin{aligned}
 \int_{\mathcal{S}} \varphi(x) d\mathbf{p}(x) &= \int_{\cup_{i=1}^n B(X_i, \varepsilon)} \varphi(x) d\mathbf{p}(x) + \int_{\cap_{i=1}^n B^c(X_i, \varepsilon)} \varphi(x) d\mathbf{p}(x) \\
 &= \int_{\cup_{i=1}^n B(X_i, \varepsilon)} \left(\varphi(x) - \varphi(X^*(x)) \right) d\mathbf{p}(x) + \sum_{i=1}^n \mathbf{p}(C_i) \varphi(X_i) \\
 &\quad + \int_{\cap_{i=1}^n B^c(X_i, \varepsilon)} \varphi(x) d\mathbf{p}(x).
 \end{aligned} \quad (3.19)$$

Since the sets C_i are disjoint, we have that

$$\sum_{i=1}^n \mathbf{p}(C_i) = \mathbf{p}(\cup_{i=1}^n C_i) = \mathbf{p}(\cup_{i=1}^n B(X_i, \varepsilon)) = 1 - \mathbf{p}(\cap_{i=1}^n B(X_i, \varepsilon)^c).$$

Using this property as well as the Lipschitz continuity and the boundedness of φ , we obtain that

$$\begin{aligned} & \left| \int_{\mathcal{X}} \varphi(x) \, d\mathbf{p}(x) - \int_{\mathcal{X}} \varphi(x) \, d\mathbf{p}^{n,\varepsilon}(x) \right| \\ & \leq \varepsilon \operatorname{Lip}(\varphi) \mathbf{p}(\cup_{i=1}^n B(X_i, \varepsilon)) + \frac{(1 - \sum_{i=1}^n \mathbf{p}(C_i))}{\sum_{i=1}^n \mathbf{p}(C_i)} \sum_{i=1}^n \mathbf{p}(C_i) |\varphi(X_i)| + \|\varphi\|_{\infty} \mathbf{p}(\cap_{i=1}^n B(X_i, \varepsilon)^c) \\ & \leq \varepsilon \operatorname{Lip}(\varphi) \mathbf{p}(\cup_{i=1}^n B(X_i, \varepsilon)) + \frac{\mathbf{p}(\cap_{i=1}^n B(X_i, \varepsilon)^c)}{\sum_{i=1}^n \mathbf{p}(C_i)} \sum_{i=1}^n \mathbf{p}(C_i) |\varphi(X_i)| + \|\varphi\|_{\infty} \mathbf{p}(\cap_{i=1}^n B(X_i, \varepsilon)^c) \\ & \leq \varepsilon \operatorname{Lip}(\varphi) \mathbf{p}(\cup_{i=1}^n B(X_i, \varepsilon)) + 2\|\varphi\|_{\infty} \mathbf{p}(\cap_{i=1}^n B(X_i, \varepsilon)^c). \end{aligned}$$

Hence, we conclude that

$$\sup_{\varphi \in \operatorname{Lip}(\mathcal{X}), \|\varphi\|_{BL} \leq 1} \left| \int_{\mathcal{X}} \varphi(x) \, d\mathbf{p}(x) - \int_{\mathcal{X}} \varphi(x) \, d\mathbf{p}^{n,\varepsilon}(x) \right| \leq \varepsilon \mathbf{p}(\cup_{i=1}^n B(X_i, \varepsilon)) + 2\mathbf{p}(\cap_{i=1}^n B(X_i, \varepsilon)^c),$$

and passing it to the expectation yields (3.18). \square

The bound in (3.18) guarantees that for every $\varepsilon > 0$, there exists an integer n_0 such that, for all $n \geq n_0$,

$$\mathbb{E}[\operatorname{d}_{BL}(\mathbf{p}, \mathbf{p}^{n,\varepsilon})] \leq 3\varepsilon.$$

Moreover, provided that \mathbf{p} is compactly supported, Lemma 3.7 provides that we can choose

$$n_0 \geq \lceil \log(\varepsilon) / \log(1 - \mathbf{p}_{\varepsilon}) \rceil.$$

Furthermore, we obtain the same result, up to a constant factor, with respect to the p -Wasserstein distance, in case \mathbf{p} is compactly supported.

Corollary 3.11. *Assume, in addition to the assumptions of Lemma 3.10, that \mathbf{p} is compactly supported. Then*

$$\mathbb{E} \left[W_p(\mathbf{p}, \mathbf{p}^{n,\varepsilon}) \right] \leq C_{\mathbf{p},\varepsilon} (\varepsilon p_{\varepsilon,n} + 2(1 - p_{\varepsilon,n})), \quad (3.20)$$

where $C_{\mathbf{p},\varepsilon}$ is given by

$$C_{\mathbf{p},\varepsilon} := 2(\operatorname{diam}(\operatorname{supp} \mathbf{p}) + 2\varepsilon)^{\frac{p-1}{p}} (\operatorname{diam}(\operatorname{supp} \mathbf{p}) + 2\varepsilon + 1).$$

Proof. Let $K := \{x \in \mathbb{R}^d : \operatorname{dist}(x, \operatorname{supp} \mathbf{p}) \leq \varepsilon\}$. Since \mathbf{p} is compactly supported, K is compact as well with

$$\operatorname{diam}(K) \leq \operatorname{diam}(\operatorname{supp} \mathbf{p}) + 2\varepsilon. \quad (3.21)$$

Observe that with probability 1 both \mathbf{p} and $\mathbf{p}^{n,\varepsilon}$ belong to $\mathcal{P}(K)$. Since K is compact, all the p -Wasserstein distances are equivalent to the 1-Wasserstein distance on $\mathcal{P}(K)$: indeed, it is easy to check that

$$W_p(\mu, \nu) \leq \operatorname{diam}(K)^{\frac{p-1}{p}} W_1(\mu, \nu) \quad \text{for every } \mu, \nu \in \mathcal{P}(K) \text{ and every } p \in [1, +\infty).$$

The Kantorovich-Rubinstein duality theorem (see, e.g., [96, Rem. 6.5]) states that

$$W_1(\mu, \nu) = \sup \left\{ \int \varphi \, d(\mu - \nu) : \varphi \in C_b(K), \operatorname{Lip}(\varphi) \leq 1 \right\} =: \operatorname{d}_{KR}(\mu, \nu) \quad (3.22)$$

for every $\mu, \nu \in \mathcal{P}(K)$; moreover, it is not difficult to see that

$$\operatorname{d}_{KR}(\mu, \nu) \leq 2(1 + \operatorname{diam}(K)) \operatorname{d}_{BL}(\mu, \nu) \quad (3.23)$$

for every $\mu, \nu \in \mathcal{P}(K)$. Finally, combining the above inequalities (3.21), (3.22), and (3.23) with the inequality (3.18) from Lemma 3.10 yields the desired bound. \square

The main issue with Lemma 3.10 and Corollary 3.11 is that the empirical measure used for estimating the integration does require the possibility of accessing and evaluating the entire measure \mathbf{p} on suitable subsets C_i of random balls. Hence, in general, it won't be practicable to use such results for estimating d_{BL} or W_p^p , unless one considers a further density estimation on such random balls. For this reason we leave it at this point and we consider a further useful result, which does not require to estimate the measure on sets. This simple result applies to a very special type of empirical functions, depending on the distance of the argument to an empirical realization of the measure \mathbf{p} .

Lemma 3.12. *Let $(\mathcal{S}, d, \mathbf{p})$ be a Polish metric-probability space and consider a function $L : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ that is continuous, bounded, and satisfies $\lim_{r \downarrow 0} L(r) = 0$. Then, for all $\delta > 0$ there exists $\varepsilon > 0$ such that*

$$\mathbb{E} \left[\int_{\mathcal{S}} L \left(\min_{i=1, \dots, n} d(x, X_i) \right) d\mathbf{p}(x) \right] \leq \delta p_{\varepsilon, n} + \|L\|_{\infty} (1 - p_{\varepsilon, n}), \quad (3.24)$$

where X_1, \dots, X_n are i.i.d. random variables on \mathcal{S} distributed according to \mathbf{p} and $n \in \mathbb{N}$.

Proof. For any $\delta > 0$ there exists $\varepsilon > 0$ such that $L(r) \leq \delta$ for any $0 \leq r \leq \varepsilon$. Let X_1, \dots, X_n be i.i.d. random variables on \mathcal{S} distributed according to \mathbf{p} and let us consider a random (ε, n) -subcovering $\mathcal{C} = \{B(X_i, \varepsilon) : i = 1, \dots, n\}$ of \mathcal{S} . We estimate

$$\begin{aligned} \int_{\mathcal{S}} L \left(\min_{i=1, \dots, n} d(x, X_i) \right) d\mathbf{p}(x) &= \int_{\cup_{i=1}^n B(X_i, \varepsilon)} L \left(\min_{i=1, \dots, n} d(x, X_i) \right) d\mathbf{p}(x) \\ &\quad + \int_{\cap_{i=1}^n B(X_i, \varepsilon)^c} L \left(\min_{i=1, \dots, n} d(x, X_i) \right) d\mathbf{p}(x) \\ &\leq \delta \mathbf{p}(\cup_{i=1}^n B(X_i, \varepsilon)) + \|L\|_{\infty} \mathbf{p}(\cap_{i=1}^n B^c(X_i, \varepsilon)). \end{aligned}$$

By taking the expectation we obtain the bound (3.24). \square

Inspired by the latter result, we come back now to the more concrete situation in which $(\mathcal{S}, d) = (\mathcal{P}(K), W_p)$ for a compact subset $K \subset \mathbb{R}^d$ and we conclude this subsection with a random version of Proposition 3.4. Its proof actually follows a similar argument as the one of Lemma 3.12.

Proposition 3.13. *Let $N \in \mathbb{N}$ and $K \subset \mathbb{R}^d$ be a compact set. Consider μ_1, \dots, μ_N i.i.d. random variables on $\mathcal{P}_p(K) = \mathcal{P}(K)$ distributed according to $\mathbf{p} \in \mathcal{P}(\mathcal{P}(K))$ and let $\vartheta \in \mathcal{P}(K)$ be fixed. Define (the random) function $F_N(\vartheta, \cdot) : \mathcal{P}(K) \rightarrow \mathbb{R}$ as in Proposition 3.4 with $(\mu_i)_{i=1}^N$ in place of $(\sigma^h)_{h \in \mathbb{N}}$. Then*

$$\mathbb{E} \left[\int_{\mathcal{P}(K)} |F_N(\vartheta, \mu) - W_p(\vartheta, \mu)|^q d\mathbf{p}(\mu) \right] \leq C(p, q, K) [(1 - p_{\varepsilon, N}) + p_{\varepsilon, N}(N^{-q} + \varepsilon^q)],$$

where $q \in [1, +\infty)$, $\varepsilon > 0$, and $C(p, q, K)$ is a constant depending only on p, q and K .

Proof. The proof is based on the one of Proposition 3.4 with some changes due to the present probabilistic setting. We rewrite the expected value above in a convenient form: for any $\delta > 0$

we have

$$\begin{aligned}
\mathbb{E} \left[\int_{\mathcal{P}(K)} |F_N(\vartheta, \mu) - W_p(\vartheta, \mu)|^q \, d\mathbf{p}(\mu) \right] &= \int_{(\mathcal{P}(K))^N} \int_{\mathcal{P}(K)} |F_N(\vartheta, \mu) - W_p(\vartheta, \mu)|^q \, d\mathbf{p}(\mu) \, d(\otimes^N \mathbf{p})(\mu_1, \dots, \mu_N) \\
&= \int_{(\mathcal{P}(K))^N} \int_{\bigcap_{i=1}^N \mathbb{B}(\mu_i, \delta)^c} |F_N(\vartheta, \mu) - W_p(\vartheta, \mu)|^q \, d\mathbf{p}(\mu) \, d(\otimes^N \mathbf{p})(\mu_1, \dots, \mu_N) \\
&\quad + \int_{(\mathcal{P}(K))^N} \int_{\bigcup_{i=1}^N \mathbb{B}(\mu_i, \delta)} |F_N(\vartheta, \mu) - W_p(\vartheta, \mu)|^q \, d\mathbf{p}(\mu) \, d(\otimes^N \mathbf{p})(\mu_1, \dots, \mu_N).
\end{aligned}$$

Recall that we can estimate the Wasserstein distance in terms of the diameter of K and that $F_N(\vartheta, \mu)$ is bounded from above by $W_p(\vartheta, \mu)$. Hence, there exists a constant $C(q, p, K)$ depending only on q, p and K such that

$$|F_N(\vartheta, \mu) - W_p(\vartheta, \mu)|^q \leq C(q, p, K) \quad \text{for every } \mu \in \mathcal{P}(K).$$

This leads to

$$\begin{aligned}
\mathbb{E} \left[\int_{\mathcal{P}(K)} |F_N(\vartheta, \mu) - W_p(\vartheta, \mu)|^q \, d\mathbf{p}(\mu) \right] &\leq C(q, p, K) \int_{(\mathcal{P}(K))^N} \mathbf{p} \left(\bigcap_{i=1}^N \mathbb{B}(\mu_i, \delta)^c \right) \, d(\otimes^N \mathbf{p})(\mu_1, \dots, \mu_N) \\
&\quad + \int_{(\mathcal{P}(K))^N} \int_{\bigcup_{i=1}^N \mathbb{B}(\mu_i, \delta)} |F_N(\vartheta, \mu) - W_p(\vartheta, \mu)|^q \, d\mathbf{p}(\mu) \, d(\otimes^N \mathbf{p})(\mu_1, \dots, \mu_N) \\
&\leq C(q, p, K)(1 - p_{\delta, N}) \\
&\quad + \int_{(\mathcal{P}(K))^N} \int_{\bigcup_{i=1}^N \mathbb{B}(\mu_i, \delta)} |F_N(\vartheta, \mu) - W_p(\vartheta, \mu)|^q \, d\mathbf{p}(\mu) \, d(\otimes^N \mathbf{p})(\mu_1, \dots, \mu_N),
\end{aligned}$$

We can assume without loss of generality that $\varepsilon \in (0, 1)$ and set $\delta := \varepsilon^p$ so that $\delta < \varepsilon$; we now bound the second integral: for fixed $\mu_1, \dots, \mu_N \in \mathcal{P}(K)$, let $\mu \in \bigcup_{i=1}^N \mathbb{B}(\mu_i, \delta)$. We can thus find $i_0 = i_0(\mu) \in \{1, \dots, N\}$ such that $W_p(\mu, \mu_{i_0}) \leq \varepsilon^p$. Proceeding along the lines of the proof of Proposition 3.4 we obtain

$$\begin{aligned}
|F_N(\vartheta, \mu) - W_p(\vartheta, \mu)| &\leq \frac{1}{N} + p^{1/p} D_{p, R}^{1/p} W_p^{1/p}(\mu, \mu_{i_0}) (1 + \mathfrak{m}_p(\mu) + \mathfrak{m}_p(\mu_{i_0}))^{1/p} \\
&\quad + W_p(\mu_{i_0}, \mu) + W_p(\mu_{1/N}, \mu) + W_p(\hat{\vartheta}_{1/N, R}, \vartheta),
\end{aligned}$$

where we have used the notation introduced before the proof of Proposition 3.4 and $R > 0$ is such that $K \subset \mathbb{B}(0, R - 1)$. Using that $W_p(\mu_{1/N}, \mu) \leq \mathfrak{m}_p(\kappa \mathcal{L}^d)/N$, see, e.g., [3, Lemma 7.1.10], and (3.9) we further obtain that

$$|F_N(\vartheta, \mu) - W_p(\vartheta, \mu)| \leq \frac{1}{N} + C(p, K, \varepsilon) \varepsilon + \varepsilon + \frac{C(p, K)}{N} \leq C'(p, K)(\varepsilon + 1/N),$$

where we have used that in (3.9) R depends only on K . This estimate doesn't depend on μ, μ_1, \dots, μ_N so that we can integrate the inequality and conclude the proof. \square

In contrast to Proposition 3.4, the latter result shows that we can construct, from a finite number of samples of the distribution \mathbf{p} , a cylinder function that approximates the Wasserstein distance in expectation. While the bound is not deterministic, we obtain a rate of convergence depending on the ε -subcovering probabilities relative to \mathbf{p} . The more such measure is locally concentrated, the tighter is the bound.

4. CONSISTENCY OF EMPIRICAL RISK MINIMIZATION

In order to extend the approximation results of Section 3.1 to more general functions than the Wasserstein distance from a reference measure, we study in this section the construction of approximants by empirical risk minimization.

We also introduce a regularizing term using the pre-Cheeger energy as additional term to cope with noisy data.

In Section 4.1, we prove a completely deterministic result, which however lacks of a convergence rate and does not offer yet a proper analysis for data corrupted by noise.

In Section 4.2 we discuss a different approach (based on [26]) which leads to an analogous kind of convergence in mean (i.e., non-deterministic) gaining an explicit order of convergence also for noisy data.

4.1. Gamma convergence with approximating measures

In this section, in order to keep things simpler, we choose the convenient exponent $p = 2$ for the order of the Wasserstein space. We consider again a compact subset $K \subset \mathbb{R}^d$ and we thus work on the space $\mathbb{W}_2 := (\mathcal{P}_2(K), W_2, \mathbf{m}) = (\mathcal{P}(K), W_2, \mathbf{m})$, where \mathbf{m} is a non-negative and finite Borel measure on $\mathcal{P}(K)$. Recall that, since K is compact, also $\mathcal{P}(K)$ is compact and W_2 metrizes the weak convergence; see, e.g., [3, Proposition 7.1.5]. We further fix a sequence $(\mathbf{m}_n)_{n \in \mathbb{N}}$ of finite and positive Borel measures on $\mathcal{P}(K)$ such that $\mathbf{m}_n \rightarrow \mathbf{m}$ (equivalently $W_2(\mathbf{m}_n, \mathbf{m}) \rightarrow 0$ as $n \rightarrow +\infty$, where, with a slight abuse of notation, we are denoting by W_2 also the Wasserstein distance on $\mathcal{P}(\mathcal{P}(K))$ with the ground metric being given by the 2-Wasserstein distance on $\mathcal{P}(K)$). We fix a Lipschitz continuous function $F : \mathcal{P}(K) \rightarrow \mathbb{R}$ and denote by $L_F \geq 0$ its Lipschitz constant w.r.t. the distance W_2 . Notice that $F \in H^{1,2}(\mathbb{W}_2)$ and that it is bounded by a constant $C_F \geq 0$. We approach the consistency problem by Γ -convergence methods [30, 8]. Let us also mention that the study of metric Sobolev spaces with changing distances/measures has also been carried out in much more general situations, see, e.g. [71, 2].

The aim of this section is to study under which conditions a suitably modified and restricted version of the Sobolev norm on $\mathbb{W}_2 = (\mathcal{P}(K), W_2, \mathbf{m})$ can be approximated in the sense of Γ -convergence by the same object defined on $(\mathcal{P}(K), W_2, \mathbf{m}_n)$. We give here the relevant definitions and we refer to Section 2 for the notion of cylinder function and related definitions.

Definition 4.1. Given a strictly decreasing and vanishing sequence $(\varepsilon_n)_n \subset (0, +\infty)$ bounding from above the sequence $W_2(\mathbf{m}_n, \mathbf{m})$, i.e.,

$$W_2(\mathbf{m}_n, \mathbf{m}) \leq \varepsilon_n \quad \text{for all } n \in \mathbb{N},$$

we define the following sequence of subsets of the cylinder functions $\mathfrak{C}(\mathcal{P}(K), C_b^1(\mathbb{R}^d))$:

$$\mathcal{V}_n := \{G := \psi \circ \mathbf{L}_\phi : \psi \in \mathcal{P}_n, \phi = (\phi_1, \dots, \phi_n), \|G\|_{CBL} \leq \eta_n\}, \quad n \in \mathbb{N},$$

where $(\phi_n)_n \subset C_c^\infty(\mathbb{R}^d)$ is a countable and dense subset of $C_c^\infty(\mathbb{R}^d)$, $\eta_n := \varepsilon_n^{-1/4}$, \mathcal{P}_n denotes the set of polynomials on \mathbb{R}^n of degree at most n , \mathbf{L}_ϕ is as in (2.10), and

$$\|G\|_{CBL} := \|G\|_\infty \vee \|DG\|_\infty \vee \text{Lip}(G, \mathcal{P}(K)) \vee \text{Lip}(DG, \mathcal{P}(K) \times K), \quad G \in \mathfrak{C}(\mathcal{P}(K), C_b^1(\mathbb{R}^d)),$$

where $\text{Lip}(\cdot, \cdot)$ is as in (2.2) and DG is defined in (2.11).

Notice that, since η_n is increasing, we have that $\mathcal{V}_n \subset \mathcal{V}_{n+1}$ and $\bigcup_n \mathcal{V}_n$ is dense in $H^{1,2}(\mathbb{W}_2)$; for the latter result we refer to [37, Proposition 4.19].

Definition 4.2. We fix $\lambda \geq 0$ and we further define the functionals

$$\begin{aligned} \mathcal{J}_n(G) &:= \begin{cases} \|G - F\|_{L^2(\mathcal{P}_2(K), \mathbf{m}_n)}^2 + \lambda \mathfrak{pCE}_{2, \mathbf{m}_n}(G), & \text{if } G \in \mathcal{V}_n, \\ +\infty & \text{if } G \in H^{1,2}(\mathbb{W}_2) \setminus \mathcal{V}_n, \end{cases} \\ \mathcal{J}(G) &:= \begin{cases} \|G - F\|_{L^2(\mathcal{P}_2(K), \mathbf{m})}^2 + \lambda \mathfrak{pCE}_{2, \mathbf{m}}(G), & \text{if } G \in \mathfrak{C}(\mathcal{P}(K), C_b^1(\mathbb{R}^d)), \\ +\infty & \text{if } G \in H^{1,2}(\mathbb{W}_2) \setminus \mathfrak{C}(\mathcal{P}(K), C_b^1(\mathbb{R}^d)), \end{cases} \\ \bar{\mathcal{J}}(G) &:= \|G - F\|_{L^2(\mathcal{P}_2(K), \mathbf{m})}^2 + \lambda \mathfrak{CE}_{2, \mathbf{m}}(G), \quad G \in H^{1,2}(\mathbb{W}_2), \end{aligned}$$

where $\mathfrak{pCE}_{2, \mathbf{m}}$ is the pre-Cheeger energy defined in (2.3) for the space $\mathbb{W}_2 = (\mathcal{P}(K), W_2, \mathbf{m})$, $\mathfrak{pCE}_{2, \mathbf{m}_n}$ is defined in the same way but with \mathbf{m}_n in place of \mathbf{m} , and $\mathfrak{CE}_{2, \mathbf{m}}$ is the Cheeger energy from (2.4) for the space \mathbb{W}_2 .

Notice that, thanks to Proposition 2.1, the pre-Cheeger energy of a cylinder function has a simple expression:

$$\mathfrak{pCE}_{2, \mathbf{m}}(G) = \int_{\mathcal{P}(K)} \int_K |DG(\mu, x)|^2 d\mu(x) d\mathbf{m}(\mu), \quad \mathfrak{pCE}_{2, \mathbf{m}_n}(G) = \int_{\mathcal{P}(K)} \int_K |DG(\mu, x)|^2 d\mu(x) d\mathbf{m}_n(\mu).$$

We start with the following preliminary result.

Lemma 4.3. *Let \mathcal{V}_n be defined as in Definition 4.1. Then, for $G \in \mathcal{V}_n$, we have that*

$$|\mathcal{J}(G) - \mathcal{J}_n(G)| \leq 2 \left((C_F + \|G\|_\infty) (\text{Lip}(G, \mathcal{P}(K)) + L_F) + 2\lambda \text{Lip}(DG, \mathcal{P}(K) \times K) \|DG\|_\infty \right) W_2(\mathbf{m}_n, \mathbf{m}).$$

Proof. The proof is only a matter of computations that we report in full for the reader's convenience.

Claim 1. *The map $A(\mu) := |G(\mu) - F(\mu)|^2$, $\mu \in \mathcal{P}(K)$, is $2(C_F + \|G\|_\infty)(\text{Lip}(G, \mathcal{P}(K)) + L_F)$ -Lipschitz continuous.*

For any $\mu, \mu' \in \mathcal{P}(K)$ we have that

$$\begin{aligned} |A(\mu) - A(\mu')| &= \left| |G(\mu) - F(\mu)|^2 - |G(\mu') - F(\mu')|^2 \right| \\ &\leq (|G(\mu)| + |G(\mu')| + |F(\mu)| + |F(\mu')|) (|G(\mu) - G(\mu')| + |F(\mu) - F(\mu')|) \\ &\leq (2C_F + 2\|G\|_\infty) (\text{Lip}(G, \mathcal{P}(K)) + L_F) W_2(\mu, \mu'). \end{aligned}$$

Claim 2. *The map $B(\mu) := \int |DG(\mu, x)|^2 d\mu(x)$, $\mu \in \mathcal{P}(K)$, is $4 \text{Lip}(DG, \mathcal{P}(K) \times K) \|DG\|_\infty$ -Lipschitz continuous.*

Let $\mu, \mu' \in \mathcal{P}_2(K)$; then

$$\begin{aligned} |B(\mu) - B(\mu')| &= \left| \int |DG(\mu, x)|^2 d\mu(x) - \int |DG(\mu', y)|^2 d\mu'(y) \right| \\ &\leq \int \left((|DG(\mu, x)| + |DG(\mu', y)|) |DG(\mu, x) - DG(\mu', y)| \right) d\gamma(x, y) \\ &\leq 2 \text{Lip}(DG) \|DG\|_\infty \int (W_2(\mu, \mu') + |x - y|) d\gamma(x, y) \\ &\leq 4 \text{Lip}(DG) \|DG\|_\infty W_2(\mu, \mu'), \end{aligned}$$

where $\gamma \in \Gamma_o(\mu, \mu')$ (cf. (2.7)). In the last inequality we used that $W_1(\mu, \mu') \leq W_2(\mu, \mu')$, which follows immediately from Jensen's inequality.

Claim 3. Conclusion.

Using the representation of the pre-Cheeger energy coming from (2.12), we find that

$$\begin{aligned}
 |\mathcal{J}(G) - \mathcal{J}_n(G)| &\leq \left| \int |G(\mu) - F(\mu)|^2 \, d\mathbf{m}_n(\mu) - \int |G(\mu') - F(\mu')|^2 \, d\mathbf{m}(\mu') \right| \\
 &\quad + \lambda \left| \int \int |DG(\mu, x)|^2 \, d\mu(x) \, d\mathbf{m}_n(\mu) - \int \int |DG(\mu', y)|^2 \, d\mu'(y) \, d\mathbf{m}(\mu') \right| \\
 &= \left| \int A \, d\mathbf{m}_n - \int A \, d\mathbf{m} \right| + \lambda \left| \int B \, d\mathbf{m}_n - \int B \, d\mathbf{m} \right| \\
 &\leq \text{Lip}(A, \mathcal{P}(K)) \int W_2(\mu, \mu') \, d\gamma_n(\mu, \mu') + \lambda \text{Lip}(B, \mathcal{P}(K)) \int W_2(\mu, \mu') \, d\gamma_n(\mu, \mu') \\
 &\leq (2(C_F + \|G\|_\infty)(\text{Lip}(G) + L_F) + 4\lambda \text{Lip}(DG)\|DG\|_\infty) W_2(\mathbf{m}_n, \mathbf{m}),
 \end{aligned}$$

where $\gamma_n \in \Gamma_o(\mathbf{m}_n, \mathbf{m})$ (cf. (2.7), with the distance ϱ being the 2-Wasserstein distance on $\mathcal{P}(K)$). \square

Theorem 4.4. *Let \mathcal{J}_n and $\bar{\mathcal{J}}$ be defined as in Definition 4.2. Then $\mathcal{J}_n \rightarrow \bar{\mathcal{J}}$ as $n \rightarrow +\infty$ in the following sense:*

(1) *for every $(G_n)_n \subset H^{1,2}(\mathbb{W}_2)$ converging weakly in $L^2(\mathcal{P}(K), \mathbf{m})$ to $G \in H^{1,2}(\mathbb{W}_2)$, it holds*

$$\liminf_n \mathcal{J}_n(G_n) \geq \bar{\mathcal{J}}(G);$$

(2) *for every $G \in H^{1,2}(\mathbb{W}_2)$ there exists a sequence $(G_n)_n \subset H^{1,2}(\mathbb{W}_2)$ converging strongly to G in $H^{1,2}(\mathbb{W}_2)$ such that*

$$\limsup_n \mathcal{J}_n(G_n) \leq \bar{\mathcal{J}}(G).$$

Proof. We prove separately the lim inf and the lim sup inequalities.

Regarding point (1), we can assume without loss of generality that $G_n \in \mathcal{V}_n$ for every $n \in \mathbb{N}$; by Lemma 4.3 and the definition of \mathcal{V}_n in Definition 4.1, we have

$$\mathcal{J}_n(G_n) \geq \mathcal{J}(G_n) - 2 \left(2\lambda\eta_n^2 + (C_F + \eta_n)(\eta_n + L_F) \right) W_2(\mathbf{m}_n, \mathbf{m}).$$

Passing to the \liminf_n and using the definition of η_n as well as the bound on $W_2(\mathbf{m}_n, \mathbf{m})$, we get that

$$\liminf_n \mathcal{J}_n(G_n) \geq \liminf_n \mathcal{J}(G_n) \geq \bar{\mathcal{J}}(G),$$

where we used for the latter inequality that $\text{pCE}_{2,\mathbf{m}}(\cdot) \geq \text{CE}_{2,\mathbf{m}}(\cdot)$ and that the Cheeger energy and the L^2 norm are lower semicontinuous w.r.t. the weak convergence in $L^2(\mathcal{P}(K), \mathbf{m})$ (these facts easily follow from the expression in (2.4)).

Concerning point (2), by the density of $\bigcup_n \mathcal{V}_n$ in $H^{1,2}(\mathbb{W}_2)$, we can find functions H_k and numbers $N_k \in \mathbb{N}$ such that

$$H_k \rightarrow G \text{ in } H^{1,2}(\mathbb{W}_2) \quad \text{and} \quad \mathcal{J}(H_k) \rightarrow \bar{\mathcal{J}}(G) \quad \text{as } k \rightarrow +\infty$$

and

$$H_k \in \mathcal{V}_{N_k} \quad \text{with} \quad N_k < N_{k+1} \quad \text{for every } k \in \mathbb{N}.$$

We can thus define $G_n := H_k$ if $N_k \leq n < N_{k+1}$ and $G_n = \mathbf{L}_{\phi_1}$ if $n < N_1$; observe that it holds

$$G_n \rightarrow G \text{ in } H^{1,2}(\mathbb{W}_2), \quad \mathcal{J}(G_n) \rightarrow \bar{\mathcal{J}}(G) \text{ as } n \rightarrow +\infty \text{ and } G_n \in \mathcal{V}_n, \text{ for every } n \in \mathbb{N}.$$

By Lemma 4.3, we have

$$\mathcal{J}_n(G_n) \leq \mathcal{J}(G_n) + 2 \left(2\lambda\eta_n^2 + (C_F + \eta_n)(\eta_n + L_F) \right) W_2(\mathbf{m}_n, \mathbf{m}).$$

Passing to the \limsup_n and using the definition of η_n coming from Definition 4.1, we get that

$$\limsup_n \mathcal{J}_n(G_n) \leq \bar{\mathcal{J}}(G);$$

this finishes the proof. \square

Corollary 4.5. *Let $G_n^* \in \mathcal{V}_n$ be the unique minimizer of \mathcal{J}_n and let $\tilde{G}_n \in \mathcal{V}_n$ be a quasi-minimizer in the sense that*

$$\mathcal{J}_n(\tilde{G}_n) \leq \mathcal{J}_n(G_n^*) + \gamma_n$$

for a vanishing non-negative sequence $(\gamma_n)_n$. Then $\tilde{G}_n \rightarrow G^*$ in $L^2(\mathcal{P}(K), \mathbf{m})$ and $\mathcal{J}_n(\tilde{G}_n) \rightarrow \bar{\mathcal{J}}(G^*)$ as $n \rightarrow +\infty$, where G^* is the unique minimizer of $\bar{\mathcal{J}}$. If $\lambda > 0$, the convergence holds in $H^{1,2}(\mathbb{W}_2)$.

Proof. The strict convexity, lower semicontinuity and coercivity of the above functionals give the existence and uniqueness of their minimizers.

It is enough to show that from any subsequence of $(\tilde{G}_n)_{n \in \mathbb{N}}$ it is possible to extract a further subsequence such that the above convergences hold. Let us consider an unrelabelled subsequence of $(\tilde{G}_n)_{n \in \mathbb{N}}$. For $n \in \mathbb{N}$ sufficiently large we have by Lemma 4.3 that

$$\begin{aligned} \mathcal{J}(\tilde{G}_n) &\leq \mathcal{J}_n(\tilde{G}_n) + 2 \left(2\lambda\eta_n^2 + (C_F + \eta_n)(\eta_n + L_F) \right) W_2(\mathbf{m}_n, \mathbf{m}) \\ &\leq \mathcal{J}_n(G_n^*) + \gamma_n + 1 \\ &\leq \mathcal{J}_n(L_{\phi_1}) + 2 \\ &< C < +\infty. \end{aligned}$$

In turn, since \mathcal{J} is coercive, the sequence $(\tilde{G}_n)_n$ is uniformly bounded in $L^2(\mathcal{P}(K), \mathbf{m})$ and thus it admits a (unrelabelled) weakly converging subsequence with limit denoted by $H \in L^2(\mathcal{P}(K), \mathbf{m})$. For any given $G \in H^{1,2}(\mathbb{W}_2)$, we can find, thanks to the convergence deduced in Theorem 4.4, a sequence $(G_n)_n$ converging to G in $H^{1,2}(\mathbb{W}_2)$ and such that $\mathcal{J}_n(G_n) \rightarrow \bar{\mathcal{J}}(G)$. We thus have

$$\bar{\mathcal{J}}(G) = \lim_n \mathcal{J}_n(G_n) \geq \liminf_n \mathcal{J}_n(G_n^*) \geq \liminf_n \mathcal{J}_n(\tilde{G}_n) - \gamma_n \geq \bar{\mathcal{J}}(H).$$

Since $G \in H^{1,2}(\mathbb{W}_2)$ was chosen arbitrarily, this shows that $H = G^*$. Moreover, choosing $G = G^*$, the above inequalities are all equalities, showing that $\mathcal{J}_n(\tilde{G}_n) \rightarrow \bar{\mathcal{J}}(G^*)$. It only remains to prove that \tilde{G}_n converges to G^* strongly in $L^2(\mathcal{P}(K), \mathbf{m})$. First of all, in light of Lemma 4.3, it is easy to see that $\bar{\mathcal{J}}(\tilde{G}_n) \rightarrow \bar{\mathcal{J}}(G^*)$. Since $H^{1,2}(\mathbb{W}_2)$ is Hilbertian (cf. Theorem 2.2), $\bar{\mathcal{J}}$ is quadratic and thus satisfies the parallelogram identity

$$\frac{1}{4}\bar{\mathcal{J}}(F_1 - F_0) = \frac{1}{2}\bar{\mathcal{J}}(F_0) + \frac{1}{2}\bar{\mathcal{J}}(F_1) - \bar{\mathcal{J}}\left(\frac{1}{2}F_0 + \frac{1}{2}F_1\right) \text{ for every } F_0, F_1 \in H^{1,2}(\mathbb{W}_2).$$

Choosing $F_0 = G^*$ and $F_1 = \tilde{G}_n$, we get

$$\limsup_n \frac{1}{4}\bar{\mathcal{J}}(\tilde{G}_n - G^*) = \frac{1}{2}\bar{\mathcal{J}}(G^*) + \limsup_n \left(\frac{1}{2}\bar{\mathcal{J}}(\tilde{G}_n) - \bar{\mathcal{J}}\left(\frac{1}{2}G^* + \frac{1}{2}\tilde{G}_n\right) \right) \leq 0,$$

where the inequality is due to the lower semicontinuity of $\bar{\mathcal{J}}$ w.r.t. the weak $L^2(\mathcal{P}(K), \mathbf{m})$ convergence. Since

$$\bar{\mathcal{J}}(\tilde{G}_n - G^*) = \|G^* - \tilde{G}_n\|_{L^2(\mathcal{P}(K), \mathbf{m})}^2 + \lambda \text{CE}_{2, \mathbf{m}}(G^* - \tilde{G}_n),$$

we have indeed verified the convergence in $L^2(\mathcal{P}(K), \mathbf{m})$ (resp. in $H^{1,2}(\mathbb{W}_2)$ if $\lambda > 0$). \square

4.2. Regularized least squares approximation of bounded functions

This section serves as probabilistic counterpart of the previous Section 4.1, doing, in a way, what Section 3.2 does compared to Section 3.1.

The final result of Section 4.1 states that we can approximate the minimizer of $\bar{\mathcal{J}}$ (basically the Sobolev norm in $H^{1,2}(\mathbb{W}_2)$ with quasi-minimizers of \mathcal{J}_n (restrictions of the Sobolev norm in $H^{1,2}(\mathcal{P}(K), W_2, \mathbf{m}_n)$); however Corollary 4.5 does not provide us with any rate of such convergence. In this section we show that, at the cost of introducing some randomness and thus obtaining only estimates in expected value, we are able to approximate any given bounded function $F : \mathcal{P}(K) \rightarrow \mathbb{R}$ with minimizers of (rescaled versions of) functionals \mathcal{J}_N , which have an analogous definition to the \mathcal{J}_n of the previous section.

Let us briefly introduce the setting of this section which is very similar to the one of Section 4.1. We fix a compact $K \subset \mathbb{R}^d$ and a probability measure \mathbf{p} on $\mathcal{P}(K) = \mathcal{P}_2(K)$. We fix two natural numbers $n, N \in \mathbb{N}_0$ and consider a *finite dimensional subspace* of cylinder functions (cf. (2.9)) $\mathcal{V}_n \subset \mathfrak{C}(\mathcal{P}(K), C_b^1(K)) \subset H^{1,2}(\mathcal{P}(K), W_2, \mathbf{p})$ such that $\dim(\mathcal{V}_n) = n$ and N i.i.d. points μ_1, \dots, μ_N distributed on $\mathcal{P}(K)$ according to \mathbf{p} ; this means that we are considering as in Section 3.2 an underlying probability space $(\Omega, \mathcal{E}, \mathbb{P})$. Accordingly we define the (random) empirical measure $\mathbf{p}_N := \frac{1}{N} \sum_{j=1}^N \delta_{\mu_j}$ and we fix $\lambda \geq 0$.

The following lemma is a simple consequence of well-known linear algebra results that we report here for the reader's convenience.

Lemma 4.6. *There exists a finite family of functions $\mathcal{B}_n := \{\ell_1, \dots, \ell_n\} \subset \mathcal{V}_n$ such that*

- (1) \mathcal{B}_n is a linear basis of \mathcal{V}_n ;
- (2) $\int_{\mathcal{P}(K)} \ell_i \ell_j d\mathbf{p} = 0$ for every $i, j \in \{1, \dots, n\}$, $i \neq j$;
- (3) $\int_{\mathcal{P}(K)} |\ell_i|^2 d\mathbf{p} = 1$ for every $i = 1, \dots, n$;
- (4) $\mathbf{pCE}_{2,\mathbf{p}}(\ell_i, \ell_j) = 0$ for every $i, j \in \{1, \dots, n\}$, $i \neq j$,

where $\mathbf{pCE}_{2,\mathbf{p}}$ is defined as in (2.3).

Proof. This is an immediate consequence of the following fact: if V is a real vector space with dimension n and $\langle \cdot, \cdot \rangle_1$ and $\langle \cdot, \cdot \rangle_2$ are two scalar products on V , then there exists a basis of V which is orthonormal w.r.t. $\langle \cdot, \cdot \rangle_1$ and orthogonal w.r.t. $\langle \cdot, \cdot \rangle_2$. Indeed, let us denote by V_i the space V endowed with $\langle \cdot, \cdot \rangle_i$, $i = 1, 2$, and consider an orthonormal basis $(b_i)_{i=1}^n$ in V_1 . Let $T : V \rightarrow V$ be defined as

$$T(x) := \sum_{i=1}^n \langle x, b_i \rangle_2 b_i.$$

It is easy to verify that $\langle T(x), y \rangle_1 = \langle x, y \rangle_2$ for every $x, y \in V$ so that

$$\langle T(x), y \rangle_1 = \langle x, y \rangle_2 = \langle x, T(y) \rangle_1$$

which means that T is selfadjoint as an operator on V_1 . Moreover $\langle T(x), x \rangle_1 = |x|_2^2 > 0$ for any $x \in V \setminus \{0\}$. This implies that T is diagonalizable; i.e., there exists an orthonormal basis e_1, \dots, e_n of V_1 such that $Te_i = \lambda_i e_i$ for some non zero $\lambda_i \in \mathbb{R}$. It is easy to check that $(e_i)_{i=1}^n$ is the sought basis. \square

Recall that the pre-Cheeger energy of a cylinder function G has a simple expression thanks to Proposition 2.1:

$$\mathbf{pCE}_{2,\mathbf{p}}(G) = \int_{\mathcal{P}(K)} \int_K |DG(\mu, x)|^2 d\mu(x) d\mathbf{p}(\mu)$$

and

$$\mathbf{pCE}_{2,\mathbf{p}_N}(G) = \int_{\mathcal{P}(K)} \int_K |DG(\mu, x)|^2 d\mu(x) d\mathbf{p}_N(\mu) = \frac{1}{N} \sum_{j=1}^N \int_K |DG(\mu_j, x)|^2 d\mu_j(x),$$

respectively.

Let us define some more objects we work with in this section.

Definition 4.7. Given a function $F \in L^2(\mathcal{P}(K), \mathfrak{p}_N)$, we define the functionals

$$\begin{aligned} \mathcal{J}_{N,F}^\lambda(G) &:= \|F - G\|_{L^2(\mathcal{P}(K), \mathfrak{p}_N)}^2 + \lambda \mathfrak{pCE}_{2, \mathfrak{p}_N}(G), & G \in \mathfrak{C}(\mathcal{P}(K), C_b^1(K)), \\ \mathcal{N}_{N,F}(G) &:= \|F - G\|_{L^2(\mathcal{P}(K), \mathfrak{p}_N)}^2, & G \in L^2(\mathcal{P}(K), \mathfrak{p}_N). \end{aligned}$$

Both functionals admit a unique minimizer when restricted to \mathcal{V}_n so that we can define

$$\begin{aligned} S_N^{\lambda, n}(F) &:= \operatorname{argmin}_{G \in \mathcal{V}_n} \mathcal{J}_{N,F}^\lambda(G), \\ P_N^n(F) &:= \operatorname{argmin}_{G \in \mathcal{V}_n} \mathcal{N}_{N,F}(G). \end{aligned}$$

Our aim is to estimate the expected value of the $L^2(\mathcal{P}(K), \mathfrak{p})$ -norm of the difference of F and $S_N^{\lambda, n}(F)$, in case F is bounded. For that purpose, we proceed extending the approach outlined in [26].

In the next Lemma we make the form of $S_N^{\lambda, n}(F)$ as in Definition 4.7 more explicit. To do that we first need to define two matrices that are crucial in our approach.

Definition 4.8. Let $\{\ell_1, \dots, \ell_n\}$ be as in Lemma 4.6. We define the matrices $L_{N,n} \in \mathbb{R}^{N \times n}$ and $D_n \in \mathbb{R}^{n \times n}$ as

$$(L_{N,n})_{ji} := \frac{1}{\sqrt{N}} \ell_i(\mu_j), \quad (D_n)_{hi} = \frac{1}{N} \sum_{j=1}^N \int_K D\ell_i(\mu_j, x) \cdot D\ell_h(\mu_j, x) d\mu_j(x).$$

The expected value of the matrices $L_{N,n}^\top L_{N,n}$ and D_n is easily computed:

$$\mathbb{E}[L_{N,n}^\top L_{N,n}] = I_n \quad \text{and} \quad \mathbb{E}[D_n] = \Gamma_n, \quad (4.1)$$

where I_n is the identity matrix of order n and $\Gamma_n \in \mathbb{R}^{n \times n}$ is a diagonal matrix with entries

$$(\Gamma_n)_{ii} = \mathfrak{pCE}_{2, \mathfrak{p}}(\ell_i), \quad i = 1, \dots, n.$$

Lemma 4.9. Let $\{\ell_1, \dots, \ell_n\}$ be as in Lemma 4.6 and let $F \in L^2(\mathcal{P}(K), \mathfrak{p}_N)$. Then $G = S_N^{\lambda, n}(F)$ (cf. Definition 4.7) if and only if

$$G = \sum_{i=1}^n \ell_i w_i^*,$$

where $w^* \in \mathbb{R}^n$ is the unique minimizer in \mathbb{R}^n of the functional $\mathcal{J}_{N, z^F}^{\lambda, n} : \mathbb{R}^n \rightarrow \mathbb{R}$ defined as

$$\mathcal{J}_{N, z^F}^{\lambda, n}(w) := |L_{N,n}(w - z^F)|_N^2 + \lambda w^\top D_n w, \quad w \in \mathbb{R}^n, \quad (4.2)$$

where $|\cdot|_N$ denotes the Euclidean norm in \mathbb{R}^N and $z^F \in \mathbb{R}^n$ is such that

$$P_N^n F = \sum_{i=1}^n \ell_i z_i^F.$$

Equivalently w^* solves

$$(L_{N,n}^\top L_{N,n} + \lambda D_n) w^* = L_{N,n}^\top L_{N,n} z^F. \quad (4.3)$$

In particular, $S_N^{\lambda, n}$ is a linear operator.

Proof. For every $G \in \mathcal{V}_n$, we can rewrite

$$\begin{aligned} \mathcal{J}_{N,F}^\lambda(G) &= \frac{1}{N} \sum_{i=1}^N |G(\mu_i) - F(\mu_i)|^2 + \frac{\lambda}{N} \sum_{i=1}^N \|\mathrm{DG}[\mu_i]\|^2 \\ &= \frac{1}{N} \sum_{i=1}^N |G(\mu_i) - P_N^n F(\mu_i)|^2 + \frac{\lambda}{N} \sum_{i=1}^N \|\mathrm{DG}[\mu_i]\|^2 + C(n, N, F) \end{aligned}$$

where $C(n, N, F)$ is a constant depending only on F , n and N , and $\|\mathrm{DG}[\cdot]\|$ is as in (2.12). Furthermore, upon defining

$$\tilde{\mathcal{J}}_{N,F}^\lambda(G) := \frac{1}{N} \sum_{i=1}^N |G(\mu_i) - P_N^n F(\mu_i)|^2 + \frac{\lambda}{N} \sum_{i=1}^N \|\mathrm{DG}[\mu_i]\|^2$$

we obtain that

$$\mathcal{J}_{N,F}^\lambda(G) = \tilde{\mathcal{J}}_{N,F}^\lambda(G) + C(n, N, F).$$

Since $\tilde{\mathcal{J}}_{N,F}^\lambda$ and $\mathcal{J}_{N,F}^\lambda$ differ only by a constant, they have the same minimizer. For a given $G \in \mathcal{V}_n$ there exists a unique $w^G \in \mathbb{R}^n$ such that

$$G = \sum_{i=1}^n \ell_i w_i^G$$

and it is immediate to check that $\tilde{\mathcal{J}}_{N,F}^\lambda(G) = \mathcal{J}_{N,z^F}^{\lambda,n}(w^G)$. This concludes the proof. \square

Remark 4.10. We note that (4.3) can equivalently be stated as

$$(L_{N,n}^\top L_{N,n} + \lambda D_n) w^* = y^F,$$

where $y^F \in \mathbb{R}^n$ with

$$y_i^F = \frac{1}{N} \sum_{j=1}^N F(\mu_j) \ell_i(\mu_j), \quad 1 = 1, \dots, n.$$

In particular, we may obtain the solution w^* by solving a linear equation.

Next result is an auxiliary large deviation bound, which allows us to control the spectral behavior of the matrices introduced in Definition 4.8.

Proposition 4.11. *Let $L_{N,n}$ and D_n be defined as in Definition 4.8 and let I_n and Γ_n be as in (4.1). Then*

$$\mathbb{P}(\|L_{N,n}^\top L_{N,n} - I_n\| > \delta) \leq 2n \exp\left\{-N \frac{c_\delta}{K_\lambda(n)}\right\} \quad \text{for every } 0 < \delta < 1 \quad (4.4)$$

and, for every $0 < \delta < 1$,

$$\mathbb{P}\left(\|L_{N,n}^\top L_{N,n} + \lambda D_n - (I_n + \lambda \Gamma_n)\| > \delta \sigma_{\max,n}^\lambda\right) \leq 2n \exp\left\{-N \sigma_{\min,n}^\lambda \frac{c_\delta}{K_\lambda(n)}\right\}, \quad (4.5)$$

where

$$\begin{aligned} K_\lambda(n) &:= \sup_{\mu \in \mathcal{P}(K)} \sum_{i=1}^n \left(|\ell_i(\mu)|^2 + \lambda \int_K |\mathrm{D}\ell_i(\mu, x)|^2 \mathrm{d}\mu(x) \right), \\ \sigma_{\min,n}^\lambda &:= 1 + \lambda \min_{i=1, \dots, n} \mathrm{pCE}_{2,p}(\ell_i), \quad \sigma_{\max,n}^\lambda := 1 + \lambda \max_{i=1, \dots, n} \mathrm{pCE}_{2,p}(\ell_i), \\ c_\delta &:= (1 + \delta) \log(1 + \delta) - \delta, \end{aligned}$$

and $\|\cdot\|$ denotes the spectral norm.

Proof. The first bound follows precisely as in the proof of [26, Theorem 1]. The second one analogously: $L_{N,n}^\top L_{N,n} + \lambda D_n$ is the sum of N symmetric and positive semi-definite matrices X_k , $k = 1, \dots, N$, which are i.i.d. copies of the matrix X with entries

$$X_{ih} := X_{ih}^1 + X_{ih}^2 := \frac{1}{N} \ell_i(\mu) \ell_h(\mu) + \frac{\lambda}{N} \int_K D\ell_i(\mu, x) \cdot D\ell_h(\mu, x) d\mu(x), \quad i, h \in \{1, \dots, n\},$$

with μ distributed according to \mathfrak{p} .

Indeed, since both X^1 and X^2 are symmetric and positive semi-definite matrices, so is X and thus its spectral norm can be bounded by its trace; i.e., we have that

$$\|X\| \leq \frac{1}{N} \sum_{i=1}^n \left(|\ell_i(\mu)|^2 + \lambda \int_K |D\ell_i(\mu, x)|^2 d\mu(x) \right) \quad \text{for a.e. } \mu \in \mathcal{P}(K).$$

Invoking the definition of $K_\lambda(n)$, this immediately yields the uniform bound

$$\|X\| \leq \frac{K_\lambda(n)}{N}.$$

In turn, for every $0 < \delta < 1$, the matrix Chernoff inequality, cf. [94, Theorem 1.1], yields that

$$\begin{aligned} \mathbb{P} \left(\lambda_{\min}(L_{N,n}^\top L_{N,n} + \lambda D_n) - \sigma_{\min,n}^\lambda \leq -\delta \sigma_{\min,n}^\lambda \right) &= \mathbb{P} \left(\lambda_{\min}(L_{N,n}^\top L_{N,n} + \lambda D_n) \leq (1 - \delta) \sigma_{\min,n}^\lambda \right) \\ &\leq n \left(\frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^{\frac{N\sigma_{\min,n}^\lambda}{K_\lambda(n)}} \\ &\leq n \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\frac{N\sigma_{\min,n}^\lambda}{K_\lambda(n)}} \end{aligned}$$

and

$$\begin{aligned} \mathbb{P} \left(\lambda_{\max}(L_{N,n}^\top L_{N,n} + \lambda D_n) - \sigma_{\max,n}^\lambda \geq \delta \sigma_{\max,n}^\lambda \right) &= \mathbb{P} \left(\lambda_{\max}(L_{N,n}^\top L_{N,n} + \lambda D_n) \geq (1 + \delta) \sigma_{\max,n}^\lambda \right) \\ &\leq n \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\frac{N\sigma_{\max,n}^\lambda}{K_\lambda(n)}} \\ &\leq n \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\frac{N\sigma_{\min,n}^\lambda}{K_\lambda(n)}} \end{aligned}$$

so that

$$\mathbb{P} \left(\|L_{N,n}^\top L_{N,n} + \lambda D_n - (I_n + \lambda \Gamma_n)\| > \delta \sigma_{\max,n}^\lambda \right) \leq 2n \exp \left\{ -N \sigma_{\min,n}^\lambda \frac{c_\delta}{K_\lambda(n)} \right\}.$$

□

The condition (4.6) below ensures that the probabilities in Proposition 4.11 are small: we think to fix the dimension n of the subspace \mathcal{V}_n and an order $r > 0$. Then, up to choosing enough point evaluations N , we obtain that the above probability is smaller than N^{-r} . The required condition reads as

$$\frac{N}{\log(N)} \geq \frac{(1+r)K_\lambda(n)}{\sigma_{\min,n}^\lambda c_{\frac{1}{2\sigma_{\max,n}^\lambda}}}, \quad N \geq 2, \quad (4.6)$$

which further implies that

$$\frac{N}{\log(N)} \geq \frac{(1+r)K_\lambda(n)}{c_{1/2}}.$$

Indeed, given (4.6) and choosing $\delta = 1/2$ in (4.4) as well as $\delta = (2\sigma_{\max,n}^\lambda)^{-1}$ in (4.5) leads to

$$\begin{aligned} & \mathbb{P} \left(\|(L_{N,n}^\top L_{N,n} + \lambda D_n) - (I_n + \lambda \Gamma_n)\| > \frac{1}{2} \right) + \mathbb{P} \left(\|L_{N,n}^\top L_{N,n} - I_n\| > \frac{1}{2} \right) \\ & \leq 4 \frac{n}{N} N^{-r} \leq 4 \frac{K_\lambda(n)}{N} N^{-r} \leq \frac{4c_{1/2}}{\log(N)(1+r)} N^{-r} \leq \frac{4c_{1/2}}{\log(2)} N^{-r} \leq N^{-r}, \end{aligned} \quad (4.7)$$

where we have used that $K_\lambda(n) \geq n$ and that $N \geq 2$.

Notice moreover that, if $\|L_{N,n}^\top L_{N,n} - I_n\| \leq \frac{1}{2}$, we get

$$\|(L_{N,n}^\top L_{N,n})^{-1}\| \leq \sum_{k=0}^{\infty} \|L_{N,n}^\top L_{N,n} - I_n\|^k \leq \frac{1}{1 - \frac{1}{2}} = 2. \quad (4.8)$$

Analogously, if $\|(L_{N,n}^\top L_{N,n} + \lambda D_n) - (I_n + \lambda \Gamma_n)\| \leq \frac{1}{2}$, we get, setting $B := (L_{N,n}^\top L_{N,n} + \lambda D_n)$ and $\Gamma_\lambda := (I_n + \lambda \Gamma_n)$, that

$$\sigma_{\min,n}^\lambda \|B \Gamma_\lambda^{-1} - I_n\| \leq \|B - \Gamma_\lambda\| \leq \frac{1}{2}$$

so that

$$\sigma_{\min,n}^\lambda \|B^{-1}\| \leq \|B^{-1} \Gamma_\lambda\| \leq \frac{1}{1 - \frac{1}{2\sigma_{\min,n}^\lambda}}.$$

Consequently, we have that

$$\|(L_{N,n}^\top L_{N,n} + \lambda D_n)^{-1}\| = \|B^{-1}\| \leq \frac{1}{\frac{1}{2} + \lambda \min_{i=1,\dots,n} \mathbf{pCE}_{2,\mathbf{p}}(\ell_i)}. \quad (4.9)$$

Now, we are ready to present the first result of this section, where the data are noiseless samples of the function F we intend to recover.

Theorem 4.12. *Let $M > 0$ and let $F : \mathcal{P}(K) \rightarrow [-M, M]$ be a measurable function and let F^* be defined as*

$$F^* := -M \vee S_N^{\lambda,n}(F) \wedge M,$$

where $S_N^{\lambda,n}$ is as in Definition 4.7; let $r > 0$ be given and assume that $n, N \in \mathbb{N}$ satisfy (4.6). Then

$$\begin{aligned} & \mathbb{E} \left[\|F - F^*\|_{L^2(\mathcal{P}(K), \mathbf{p})}^2 \right] \\ & \leq 2e(F, n) \left(1 + \frac{c_{1/2}}{\log(N)(1+r)(1/2 + \lambda \mu_{\min,n})^2} \right) + 4\lambda \mathbf{pCE}_{2,\mathbf{p}\mathbf{N}}(P_n F) + 2M^2 N^{-r}, \end{aligned}$$

where $P_n F$ is the $L^2(\mathcal{P}(K), \mathbf{p})$ -orthogonal projection of F onto \mathcal{V}_n and

$$\begin{aligned} e(F, n) & := \|P_n F - F\|_{L^2(\mathcal{P}(K), \mathbf{p})}^2, \\ \mu_{\min,n} & := \min_{i=1,\dots,n} \mathbf{pCE}_{2,\mathbf{p}}(\ell_i). \end{aligned}$$

Proof. Recall that we have fixed an underlying probability space $(\Omega, \mathcal{E}, \mathbb{P})$; we consider a partition of Ω into the two sets

$$\Omega_+ := \left\{ \omega \in \Omega : \|(L_{N,n}^\top L_{N,n} + \lambda D_n) - (I_n + \lambda \Gamma_n)\| \leq \frac{1}{2} \wedge \|L_{N,n}^\top L_{N,n} - I_n\| \leq \frac{1}{2} \right\}$$

and its complement $\Omega_- := \Omega \setminus \Omega_+$. Notice that (4.7) gives $\mathbb{P}(\Omega_-) \leq N^{-r}$. We have then

$$\mathbb{E} \left[\|F - F^*\|_{L^2(\mathcal{P}(\Omega), \mathbf{p})}^2 \right] \leq \int_{\Omega_+} \|F - F^*\|_{L^2(\mathcal{P}(\Omega), \mathbf{p})}^2 d\mathbb{P} + 2M^2 N^{-r}.$$

We are left to estimate

$$\int_{\Omega_+} \|F - F^*\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 d\mathbb{P} \leq \int_{\Omega_+} \|F - S_N^{\lambda, n}(F)\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 d\mathbb{P}.$$

Since by Lemma 4.9 the operator $S_N^{\lambda, n}$ is linear, we have

$$F - S_N^{\lambda, n}(F) = (I - S_N^{\lambda, n})H + P_n F - S_N^{\lambda, n}(P_n F),$$

where $H = F - P_n F$.

Observe that for $\xi \in \mathcal{V}_n$ and a draw in Ω_+ it holds

$$\int_{\mathcal{P}(K)} |\xi|^2 d\mathfrak{p} \leq 2 \int_{\mathcal{P}(K)} |\xi|^2 d\mathfrak{p}_N. \quad (4.10)$$

Indeed, there are coefficients $\theta_i \in \mathbb{R}$, $i = 1, \dots, n$, such that $\xi = \sum_{i=1}^n \theta_i \ell_i$, and in turn

$$\int_{\mathcal{P}(K)} |\xi|^2 d\mathfrak{p} = \sum_{i=1}^n |\theta_i|^2 = |\theta|_n^2$$

as well as

$$\int_{\mathcal{P}(K)} |\xi|^2 d\mathfrak{p}_N = \frac{1}{N} \sum_{j=1}^N |\xi(\mu_j)|^2 = \frac{1}{N} \sum_{j=1}^N \sum_{i=1}^n \sum_{h=1}^n \theta_i \theta_h \ell_i(\mu_j) \ell_h(\mu_j) = L_{N,n}^\top L_{N,n} \theta \cdot \theta \geq \frac{1}{2} |\theta|_n^2,$$

where we have used (4.8) and denoted by $|\cdot|_n$ the Euclidean norm in \mathbb{R}^n . By the very definition of $S_N^{\lambda, n}$ we further find that

$$\begin{aligned} \|S_N^{\lambda, n}(G) - G\|_{L^2(\mathcal{P}(K), \mathfrak{p}_N)}^2 &\leq \|S_N^{\lambda, n}(G) - G\|_{L^2(\mathcal{P}(K), \mathfrak{p}_N)}^2 + \lambda \mathfrak{p} \mathbf{CE}_{2, \mathfrak{p}_N}(S_N^{\lambda, n}(G)) \\ &\leq \|G - G\|_{L^2(\mathcal{P}(K), \mathfrak{p}_N)}^2 + \lambda \mathfrak{p} \mathbf{CE}_{2, \mathfrak{p}_N}(G) \\ &= \lambda \mathfrak{p} \mathbf{CE}_{2, \mathfrak{p}_N}(G) \end{aligned}$$

for any $G \in \mathcal{V}_n$. Combining this inequality with (4.10) yields

$$\begin{aligned} \int_{\Omega_+} \|P_n F - S_N^{\lambda, n}(P_n F)\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 d\mathbb{P} &\leq 2 \int_{\Omega_+} \|P_n F - S_N^{\lambda, n}(P_n F)\|_{L^2(\mathcal{P}(K), \mathfrak{p}_N)}^2 d\mathbb{P} \\ &\leq 2\lambda \mathfrak{p} \mathbf{CE}_{2, \mathfrak{p}_N}(P_n F). \end{aligned}$$

It only remains to estimate

$$\int_{\Omega_+} \|(I - S_N^{\lambda, n})H\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 d\mathbb{P}.$$

This follows precisely as in [26]: since H is $L^2(\mathcal{P}(K), \mathfrak{p})$ -orthogonal to \mathcal{V}_n and $S_N^{\lambda, n} H$ belongs to \mathcal{V}_n , we have that

$$\|(I - S_N^{\lambda, n})H\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 = \|H\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 + \|S_N^{\lambda, n} H\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2$$

and

$$\|S_N^{\lambda, n} H\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 = \sum_{i=1}^n |a_i|^2$$

where $a = (a_i)_{i=1}^n \in \mathbb{R}^n$ is the solution of

$$(L_{N,n}^\top L_{N,n} + \lambda D_n) a = y^H$$

with $y^H \in \mathbb{R}^n$ defined as $y_i^H := \int_{\mathcal{P}(K)} H(\mu) \ell_i(\mu) d\mathfrak{p}_N(\mu)$, $i = 1, \dots, n$; cf. Remark 4.10. Using (4.9), we can obtain as in [26] that

$$\int_{\Omega_+} \|S_N^{\lambda, n} H\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 d\mathbb{P} \leq \alpha_{\lambda, n}^2 \frac{K_\lambda(n)}{N} \|H\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2,$$

where

$$\alpha_{\lambda,n} := \frac{1}{\frac{1}{2} + \lambda \min_{i=1,\dots,n} \mathfrak{pCE}_{2,\mathfrak{p}}(\ell_i)}. \quad (4.11)$$

Putting everything together we get

$$\begin{aligned} \mathbb{E} \left[\|F - F^*\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 \right] &\leq \int_{\Omega_+} \|F - F^*\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 d\mathbb{P} + 2M^2 N^{-r} \\ &\leq 2 \int_{\Omega_+} \|(I - S_N^{\lambda,n})H\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 d\mathbb{P} + 4\lambda \mathfrak{pCE}_{2,\mathfrak{p}_N}(P_n F) + 2M^2 N^{-r} \\ &\leq 2\|H\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 + 2\alpha_{\lambda,n}^2 \frac{K_\lambda(n)}{N} \|H\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 + 4\lambda \mathfrak{pCE}_{2,\mathfrak{p}_N}(P_n F) + 2M^2 N^{-r}. \end{aligned}$$

Finally, recalling the definition of H and employing (4.6) yield that

$$\begin{aligned} \mathbb{E} \left[\|F - F^*\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 \right] &\leq \|P_n F - F\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 \left(2 + \frac{2\alpha_{\lambda,n}^2 c_{1/2}}{\log(N)(1+r)} \right) \\ &\quad + 4\lambda \mathfrak{pCE}_{2,\mathfrak{p}_N}(P_n F) + 2M^2 N^{-r}. \end{aligned}$$

□

We assume now to know the target function F at the (random) points μ_1, \dots, μ_N only up to some noise. In particular, we assume that we have at our disposal the noisy values

$$\tilde{F}(\mu_j) := F(\mu_j) + \eta_j, \quad j = 1, \dots, N, \quad (4.12)$$

where $(\eta_j)_{j=1}^N$ are i.i.d. random variables with variance bounded by $\sigma^2 > 0$. Notice that the (random) function \tilde{F} belongs to $L^2(\mathcal{P}(K), \mathfrak{p}_N)$.

Theorem 4.13. *Let $M > 0$, let $F : \mathcal{P}(K) \rightarrow [-M, M]$ be a measurable function and let \tilde{F}^* be defined as*

$$\tilde{F}^* := -M \vee S_N^{\lambda,n}(\tilde{F}) \wedge M,$$

where $S_N^{\lambda,n}$ is as in Definition 4.7 and \tilde{F} is as in (4.12) based on the noisy data $(y_i)_{i=1}^N$ whose i.i.d. errors $(\eta_i)_{i=1}^N$ have variance bounded by $\sigma^2 > 0$. Let $r > 0$ be given and assume that $n, N \in \mathbb{N}$ satisfy (4.6). Then

$$\begin{aligned} \mathbb{E} \left[\|F - \tilde{F}^*\|_{L^2(\mathcal{P}(K), \mathfrak{p})}^2 \right] &\leq 2e(F, n) \left(1 + \frac{c_{1/2}}{\log(N)(1+r)(1/2 + \lambda\mu_{\min,n})^2} \right) + 8\lambda \mathfrak{pCE}_{2,\mathfrak{p}_N}(P_n F) + \\ &\quad + 4 \frac{\sigma^2}{(1 + \lambda\mu_{\min,n})^2} \frac{n}{N} + 2M^2 N^{-r}, \end{aligned}$$

where $e(F, n)$ and $\mu_{\min,n}$ are as in Theorem 4.12.

Proof. The proof follows the one of Theorem 4.12 and it is also inspired by [26, Theorem 3]. We split again Ω into the sets Ω_+ and Ω_- . Performing the same calculations as before, we find that

$$\mathbb{E} \left[\|F - \tilde{F}^*\|_{L^2(\mathcal{P}(\Omega), \mathfrak{p})}^2 \right] \leq \int_{\Omega_+} \|F - S_N^{\lambda,n}(\tilde{F})\|_{L^2(\mathcal{P}(\Omega), \mathfrak{p})}^2 d\mathbb{P} + 2M^2 N^{-r}.$$

We can write

$$F - S_N^{\lambda,n}(\tilde{F}) = (I - S_N^{\lambda,n})(H + P_n F) - S_N^{\lambda,n}(\eta)$$

where $H = F - P_n F$, P_n is the orthogonal $L^2(\mathcal{P}(K), \mathfrak{p})$ projection on \mathcal{V}_n and $\eta \in L^2(\mathcal{P}(K), \mathfrak{m}_N)$ is the (random) function taking values $\eta(\mu_j) := \eta_j$, $j = 1, \dots, N$.

Following precisely the same steps as in the previous proof, we obtain the bound

$$\begin{aligned} \int_{\Omega_+} \|F - S_N^{\lambda,n}(\tilde{F})\|_{L^2(\mathcal{P}(\Omega),\mathfrak{p})}^2 d\mathbb{P} &\leq 2\|H\|_{L^2(\mathcal{P}(K),\mathfrak{p})}^2 + 8\lambda\mathfrak{p}\mathbf{CE}_{2,\mathfrak{p}_N}(P_n F) \\ &\quad + \int_{\Omega_+} \left(2\|S_N^{\lambda,n}(H)\|_{L^2(\mathcal{P}(K),\mathfrak{p})}^2 + 4\|S_N^{\lambda,n}(\eta)\|_{L^2(\mathcal{P}(K),\mathfrak{p})}^2 \right) d\mathbb{P}. \end{aligned}$$

As before, we have that

$$\|S_N^{\lambda,n}(H)\|_{L^2(\mathcal{P}(K),\mathfrak{p})}^2 = \sum_{i=1}^n |a_i|^2 \quad \text{and} \quad \|S_N^{\lambda,n}(\eta)\|_{L^2(\mathcal{P}(K),\mathfrak{p})}^2 = \sum_{i=1}^n |v_i|^2$$

where the vectors $a, v \in \mathbb{R}^n$ are the solutions of

$$(L_{N,n}^\top L_{N,n} + \lambda D_n)a = y^H \quad \text{and} \quad (L_{N,n}^\top L_{N,n} + \lambda D_n)v = y^\eta,$$

with $y_i^H = N^{-1} \sum_{j=1}^N \ell_i(\mu_j) H(\mu_j)$ and $y_i^\eta = N^{-1} \sum_{j=1}^N \ell_i(\mu_j) \eta_j$ for $i = 1, \dots, n$. Repeating the computations of [26, Theorem 3], using also (4.8) and (4.9), then leads to

$$\int_{\Omega_+} \left(2\|S_N^{\lambda,n}(H)\|_{L^2(\mathcal{P}(K),\mathfrak{p})}^2 + 4\|S_N^{\lambda,n}(\eta)\|_{L^2(\mathcal{P}(K),\mathfrak{p})}^2 \right) d\mathbb{P} \leq 2\alpha_{\lambda,n}^2 \frac{K_\lambda(n)}{N} \|H\|_{L^2(\mathcal{P}(K),\mathfrak{p})}^2 + 4\alpha_{\lambda,n}^2 \sigma^2 \frac{n}{N},$$

where $\alpha_{\lambda,n}$ is as in (4.11). Combining all the estimates and doing some simple manipulations finally lead to the claimed upper bound. \square

Remark 4.14. In practical applications it is, however, not always simple to determine the best approximating finite dimensional subspace $\mathcal{V}_n \subset \mathfrak{E}(\mathcal{P}(K), \mathbf{C}_b^1(K))$. In particular, optimizing the best approximation error $e(F, n) = \|P_n F - F\|_{L^2(\mathcal{P}(K),\mathfrak{p})}^2$ over \mathcal{V}_n remains a difficult problem. For this reason, a possible heuristic remedy by employment of neural networks is outlined in Section 6.3 below.

Remark 4.15. The results in this section can most generally be cast within the context of infinitesimally Hilbertian metric measure spaces, denoted as $(\mathcal{S}, \mathbf{d}, \mathbf{m})$. These spaces are defined as abstract metric spaces where the function space $H^{1,2}(\mathcal{S}, \mathbf{d}, \mathbf{m})$ forms a Hilbert space or, equivalently, where the associated 2-Cheeger energy exhibits quadratic behavior, as elucidated, for instance, in [41, Section 4.3].

By formulating the Tikhonov approximation for bounded functions mapping from \mathcal{S} to the real numbers, and by replacing the pre-Cheeger energy $\mathfrak{p}\mathbf{CE}_{2,\mathfrak{p}_N}$ with the abstract Cheeger energy $\mathbf{CE}_{2,\mathfrak{p}_N}$ as presented in equation (2.5), Theorem 4.12 and Theorem 4.13 can be straightforwardly extended to accommodate suitably chosen abstract sequence of finite-dimensional subspaces $\mathcal{V}_n \subset \text{Lip}_b(\mathcal{S}, \mathbf{d})$.

Consequently, the fundamental merit of the results in this paper can be understood in their ability to specify such an abstract framework [36] to the interesting case of $\mathcal{S} = \mathcal{P}(K)$. In fact, in this particular scenario, the results are practically and computationally applicable due to the density of cylinder functions, their numerical approximability, and the explicit computability of the pre-Cheeger energy $\mathfrak{p}\mathbf{CE}_{2,\mathfrak{p}_N}$.

5. SOLVING EULER-LAGRANGE EQUATIONS

We consider once again the functional $\bar{\mathcal{J}} : H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m}) \rightarrow \mathbb{R}$ from Section 4.1, but rather study its Euler-Lagrange equation than directly its minimization. We recall that

$$\bar{\mathcal{J}}(G) = \left\| G - \tilde{F} \right\|_{L_{\mathbf{m}}^2}^2 + \lambda \mathbf{CE}_{2,\mathbf{m}}(G), \quad G \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m}), \quad (5.1)$$

where $\|\cdot\|_{L_{\mathbf{m}}^2} := \|\cdot\|_{L^2(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m})}$ and $\tilde{F} = F + \eta$ for some target function $F \in L^2(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m})$ and some unknown additive noise $\eta \in L^2(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m})$.

Theorem 5.1 (Euler-Lagrange equation). *The functional $\bar{\mathcal{J}}$ from (5.1) has a unique minimizer, which is equivalently the unique solution of the linear equation*

$$(G - \tilde{F}, H)_{L_m^2} + \lambda \text{CE}_{2,m}(G, H) = 0 \quad \text{for all } H \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m}), \quad (5.2)$$

where $\text{CE}_{2,m}(\cdot, \cdot)$ is introduced in Remark 2.3.

Proof. The proof follows from the direct method of calculus of variations. Indeed, we have that the functional $\bar{\mathcal{J}}$, cf. (5.1), is Fréchet-differentiable, strictly convex, and weakly coercive, and thus has a unique minimizer; we refer to [100, Theorem 25.E]. Furthermore, by [100, Theorem 25.F], the unique minimizer of $\bar{\mathcal{J}}$ is equivalently the unique solution of the Euler–Lagrange equation

$$G \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m}) : \quad \langle \delta \bar{\mathcal{J}}(G), H \rangle = 0 \quad \text{for all } H \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m}), \quad (5.3)$$

where $\delta \bar{\mathcal{J}}$ denotes the Fréchet-derivative of $\bar{\mathcal{J}}$ and $\langle \cdot, \cdot \rangle$ signifies the duality pairing between $H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m})$ and its dual space. A straightforward calculation reveals that

$$\frac{1}{2} \langle \delta \bar{\mathcal{J}}(G), H \rangle = (G - \tilde{F}, H)_{L_m^2} + \lambda \text{CE}_{2,m}(G, H), \quad G, H \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathbf{m}), \quad (5.4)$$

and thus concludes the proof. \square

In particular, rather than minimizing the functional (5.1), one may solve the linear equation (5.2).

We shall now make a link to Section 4.2. In particular, let K be a compact subset of \mathbb{R}^d and \mathbf{p} a probability measure on $\mathcal{P}(K)$. We further consider a finite dimensional subspace of cylinder functions $\mathcal{V}_n \subset \mathfrak{C}(\mathcal{P}(K), C_b^1(K)) \subset H^{1,2}(\mathcal{P}(K), W_2, \mathbf{p})$ and N i.i.d. points μ_1, \dots, μ_N distributed on $\mathcal{P}(K)$ according to \mathbf{p} . As in the previous section, we define the (random) empirical measure $\mathbf{p}_N := \frac{1}{N} \sum_{i=1}^N \delta_{\mu_i}$. Then, we are interested in the minimization problem

$$\operatorname{argmin}_{G \in \mathcal{V}_n} \mathcal{J}_{N, \tilde{F}}^\lambda(G), \quad (5.5)$$

where, for $G \in \mathcal{V}_N$,

$$\mathcal{J}_{N, \tilde{F}}^\lambda(G) := \left\| G - \tilde{F} \right\|_{L_{\mathbf{p}_N}^2}^2 + \lambda \mathbf{p} \text{CE}_{2, \mathbf{p}_N}(G); \quad (5.6)$$

cp. Definition 4.7. Analogously as before, the ensuing result holds.

Theorem 5.2. *The optimization problem (5.5) has a unique solution, which is equivalently the unique solution of the finite dimensional linear equation*

$$(G - \tilde{F}, H)_{L_{\mathbf{p}_N}^2} + \lambda \mathbf{p} \text{CE}_{2, \mathbf{p}_N}(G, H) = 0 \quad \text{for all } H \in \mathcal{V}_n, \quad (5.7)$$

where, for $G, H \in \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$,

$$\mathbf{p} \text{CE}_{2,m}(G, H) := \int_{\mathcal{P}_2(\mathbb{R}^d)} \int_{\mathbb{R}^d} DG(\mu, x) \cdot DH(\mu, x) \, d\mu(x) \, d\mathbf{m}(\mu); \quad (5.8)$$

here, $DG(\mu, x)$ and $DH(\mu, x)$ are defined as in (2.11).

Indeed, the unique minimizer is given by $S_N^{\lambda, n}(\tilde{F})$ stemming from Definition 4.7, and thus Theorem 4.13 directly applies, which allows to explain the behavior of the solution as $N \rightarrow \infty$.

5.1. Corresponding saddle point problem

Upon introducing the operator norm

$$\|\delta\bar{\mathcal{J}}(G)\|_{\text{op}} := \sup_{\substack{H \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_{2,m}) \\ H \neq 0}} \frac{\frac{1}{2} \langle \delta\bar{\mathcal{J}}(G), H \rangle}{|H|_{H_m^{1,2}}}, \quad (5.9)$$

the equation (5.2) amounts to finding an element $G \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_{2,m})$ such that

$$\|\delta\bar{\mathcal{J}}(G)\|_{\text{op}} = 0. \quad (5.10)$$

Indeed, in light of (5.4) we have that

$$\sup_{\substack{H \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_{2,m}) \\ H \neq 0}} \frac{(G - \tilde{F}, H)_{L_m^2} + \lambda \text{CE}_{2,m}(G, H)}{|H|_{H_m^{1,2}}} \geq 0 \quad \text{for all } G \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_{2,m}),$$

with equality if and only if G is the unique minimizer of the functional $\bar{\mathcal{J}}$, or equivalently the unique solution of the Euler–Lagrange equation (5.2). For that reason, and in view of a computational framework that is outlined in Section 6.3, we are interested in the saddle point problem

$$\inf_{G \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_{2,m})} \sup_{\substack{H \in H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_{2,m}) \\ H \neq 0}} \frac{(G - \tilde{F}, H)_{L_m^2} + \lambda \text{CE}_{2,m}(G, H)}{|H|_{H_m^{1,2}}}. \quad (5.11)$$

Even for two explicitly given cylinder functions $G, H \in \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$ we won't be able to determine the Cheeger inner product $\text{CE}_{2,m}(G, H)$ and the norm $|H|_{H_m^{1,2}}$, respectively. However, thanks to our Proposition 5.3 below, we may replace those objects by the explicit pre-Cheeger inner product $\text{pCE}_{2,m}(G, H)$ from (5.8) and the norm $|H|_{pH_m^{1,2}}$ from (1.8), respectively.

Proposition 5.3. *The saddle-point problem (5.11) is equivalent to*

$$\inf_{G \in \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))} \sup_{\substack{H \in \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d)) \\ H \neq 0}} \frac{(G - \tilde{F}, H)_{L_m^2} + \lambda \text{pCE}_{2,m}(G, H)}{|H|_{pH_m^{1,2}}}. \quad (5.12)$$

Proof. For the ease of notation, let us define $\mathcal{C} := \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$, $\mathcal{H} := H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_{2,m})$ and $\mathcal{C}_0 := \mathcal{C} \setminus \{0\}$, $\mathcal{H}_0 := \mathcal{H} \setminus \{0\}$. Moreover, for $(G, H) \in \mathcal{H} \times \mathcal{H}_0$, define

$$\mathcal{F}(G, H) := \frac{(G - \tilde{F}, H)_{L_m^2} + \lambda \text{CE}_{2,m}(G, H)}{|H|_{H_m^{1,2}}}$$

and, for $(G, H) \in \mathcal{C} \times \mathcal{C}_0$,

$$\tilde{\mathcal{F}}(G, H) := \frac{(G - \tilde{F}, H)_{L_m^2} + \lambda \text{pCE}_{2,m}(G, H)}{|H|_{pH_m^{1,2}}}.$$

First of all we show that

$$\inf_{G \in \mathcal{C}} \sup_{H \in \mathcal{H}_0} \mathcal{F}(G, H) \leq \inf_{G \in \mathcal{C}} \sup_{H \in \mathcal{C}_0} \tilde{\mathcal{F}}(G, H).$$

For any $G \in \mathcal{C}$ and $\bar{H} \in \mathcal{H}_0$ we note that

$$\mathcal{F}(G, \bar{H}) = \frac{(G - \tilde{F}, \bar{H})_{L_m^2} + \lambda \int D_m G \cdot D_m \bar{H}}{|\bar{H}|_{H_m^{1,2}}} = \frac{(G - \tilde{F}, \bar{H})_{L_m^2} + \lambda \int DG \cdot D_m \bar{H}}{|\bar{H}|_{H_m^{1,2}}},$$

where $D_m \bar{H}$ and $D_m G$ are defined as in Remark 2.3 and the latter equality is due to (2.13). We further know that there exists a sequence $(H_n)_n \subset \mathcal{C}_0$ such that

$$H_n \rightarrow \bar{H} \quad \text{and} \quad DH_n \rightarrow D_m \bar{H}.$$

In turn, together with the above calculation, we obtain that

$$\mathcal{F}(G, \bar{H}) = \lim_{n \rightarrow \infty} \tilde{\mathcal{F}}(G, H_n) \leq \sup_{H \in \mathcal{C}_0} \tilde{\mathcal{F}}(G, H).$$

Since $\bar{H} \in \mathcal{H}_0$ was arbitrary, we can pass to the supremum in \mathcal{H}_0 obtaining that

$$\sup_{H \in \mathcal{H}_0} \mathcal{F}(G, H) \leq \sup_{H \in \mathcal{C}_0} \tilde{\mathcal{F}}(G, H)$$

and consequently

$$\inf_{G \in \mathcal{H}} \sup_{H \in \mathcal{H}_0} \mathcal{F}(G, H) \leq \inf_{G \in \mathcal{C}} \sup_{H \in \mathcal{H}_0} \mathcal{F}(G, H) \leq \inf_{G \in \mathcal{C}} \sup_{H \in \mathcal{C}_0} \tilde{\mathcal{F}}(G, H).$$

We now show the converse inequality: for any $G \in \mathcal{H}$ we can find an element $G_\varepsilon \in \mathcal{C}$ such that

$$|G - G_\varepsilon|_{L_m^2} < \varepsilon \quad \text{and} \quad |D_m G - DG_\varepsilon|_{L_m^2} < \varepsilon.$$

Therefore we have for any $H \in \mathcal{C}_0$ that

$$\begin{aligned} & \left| \frac{(G - \tilde{F}, H)_{L_m^2} + \lambda \int D_m G \cdot DH}{|H|_{H_m^{1,2}}} - \frac{(G_\varepsilon - \tilde{F}, H)_{L_m^2} + \lambda \int DG_\varepsilon \cdot DH}{|H|_{H_m^{1,2}}} \right| \\ & \leq \varepsilon \frac{|H|_{L_m^2} + \lambda |DH|_{L_m^2}}{|H|_{H_m^{1,2}}} \leq \varepsilon (\lambda^2 + 1)^{1/2}. \end{aligned}$$

Consequently, and in light of (2.13), we find that

$$\begin{aligned} \sup_{H \in \mathcal{C}_0} \mathcal{F}(G, H) & \geq \sup_{H \in \mathcal{C}_0} \frac{(G_\varepsilon - \tilde{F}, H)_{L_m^2} + \lambda \int DG_\varepsilon \cdot DH}{|H|_{H_m^{1,2}}} - \varepsilon (\lambda^2 + 1)^{1/2} \\ & \geq \sup_{H \in \mathcal{C}_0} \frac{(G_\varepsilon - \tilde{F}, H)_{L_m^2} + \lambda \int DG_\varepsilon \cdot DH}{|H|_{pH_m^{1,2}}} - \varepsilon (\lambda^2 + 1)^{1/2} \\ & = \sup_{H \in \mathcal{C}_0} \tilde{\mathcal{F}}(G_\varepsilon, H) - \varepsilon (\lambda^2 + 1)^{1/2} \\ & \geq \inf_{G \in \mathcal{C}} \sup_{H \in \mathcal{C}_0} \tilde{\mathcal{F}}(G, H) - \varepsilon (\lambda^2 + 1)^{1/2}. \end{aligned}$$

Passing to the limit as $\varepsilon \downarrow 0$ we thus get

$$\sup_{H \in \mathcal{H}_0} \mathcal{F}(G, H) \geq \sup_{H \in \mathcal{C}_0} \mathcal{F}(G, H) \geq \inf_{G \in \mathcal{C}} \sup_{H \in \mathcal{C}_0} \tilde{\mathcal{F}}(G, H)$$

for any $G \in \mathcal{H}$. Passing to the infimum w.r.t. $G \in \mathcal{H}$ leads to the conclusion. \square

Remark 5.4. As previously stated, in practice we neither have \mathfrak{m} nor \tilde{F} at our disposal, but only an empirical approximation

$$\frac{1}{N} \sum_{j=1}^N \delta_{\mu_j} =: \mathfrak{p}_N$$

and the corresponding (noisy) function values

$$y_j = \tilde{F}(\mu_j), \quad j \in \{1, 2, \dots, N\};$$

in particular, here we assume that $\mathfrak{m} = \mathfrak{p}$ is a probability measure. Furthermore, for the sake of actual numerical implementation, we shall again consider a finite dimensional subspace of

cylinder functions $\mathcal{V}_n \subset \mathfrak{C}(\mathcal{P}(K), C_b^1(K))$. Then, the corresponding finite dimensional empirical saddle point problem reads as

$$\inf_{G \in \mathcal{V}_n} \sup_{\substack{H \in \mathcal{V}_n \\ H \neq 0}} \frac{(G - \tilde{F}, H)_{L_{\mathbb{P}_N}^2} + \lambda \text{pCE}_{2, \mathbb{P}_N}(G, H)}{|H|_{pH_{\mathbb{P}_N}^{1,2}}}, \quad (5.13)$$

where

$$\begin{aligned} (G - \tilde{F}, H)_{L_{\mathbb{P}_N}^2} &:= \frac{1}{N} \sum_{j=1}^N (G(\mu_j) - \tilde{F}(\mu_j)) H(\mu_j) = \frac{1}{N} \sum_{j=1}^N (G(\mu_j) - y_j) H(\mu_j), \\ \text{pCE}_{2, \mathbb{P}_N}(G, H) &:= \frac{1}{N} \sum_{j=1}^N \int_{\mathbb{R}^d} DG(x, \mu_j) \cdot DH(x, \mu_j) d\mu_j(x), \end{aligned}$$

and

$$|H|_{pH_{\mathbb{P}_N}^{1,2}}^2 := \frac{1}{N} \sum_{j=1}^N \left(|H(\mu_j)|^2 + \int_{\mathbb{R}^d} |DH(x, \mu_j)|^2 d\mu_j(x) \right). \quad (5.14)$$

We emphasize again that in practice it is not simple to find the best approximating subspace \mathcal{V}_n , see Remark 4.14. The heuristic use of trained neural networks is again the computational remedy to leverage the saddle point problem (5.13) as shown in Section 6.3 below.

6. NUMERICAL EXPERIMENTS

In this section, we run some numerical tests to experimentally investigate the efficacy of some of our analytical results; in all our experiments below, the goal is to approximate the Wasserstein distance function $F_{\vartheta}^{W_2^2}$ from (1.4). For that purpose, we consider two distinct datasets.

MNIST. The MNIST dataset is a large collection of handwritten digits. Thereby, the training set $\mathcal{T}_{\text{train}}$ consists of 60 000 elements, $\{\mu_k\}_{k=1}^{60\,000}$, whereas the test set $\mathcal{T}_{\text{test}}$ is composed of 10 000 datapoints, $\{\nu_k\}_{k=1}^{10\,000}$. Each image in the MNIST dataset is given by 28×28 pixels in grayscale. For our experiments, we have normalized each image such that they correspond to probability measures. Moreover, our reference measure ϑ is the one corresponding to the barycenter of all the images of the digit "0" in the training set $\mathcal{T}_{\text{train}}$; this image has been computed with the Python open source library *POT*, cf. [35]. In the following, we use the notions of an image and its corresponding probability measure interchangeably. We have further used the *ot.emd* solver from the *POT* library to compute the 2-Wasserstein distance of each image in the training and test sets to the reference image. For the training set, we have also computed the corresponding Kantorovich potentials.

CIFAR-10. The CIFAR-10 dataset, cf. [63], is a collection of 32×32 colour images in 10 classes, whereby the training set $\mathcal{T}_{\text{train}}$ consists of 50 000 elements and the test set $\mathcal{T}_{\text{test}}$ is composed of 10 000 elements. As for the MNIST dataset, we have normalized each image so that we can work with probability measures. The reference measure ϑ was randomly picked from the training set, and happened to be the image of a deer. As before, we computed the 2-Wasserstein distances as well as the Kantorovich potentials with the *ot.emd* solver from the Python open source library *POT*.

Remark 6.1. We emphasize that the advantage of the approach presented herein is the superior evaluation time. To underline this property, we compare the time needed to determine the distance of every element in the test set to the reference measure in the case of the CIFAR-10

dataset by our approximating function from Experiment 6.6 below with two built-in Python functions from the *POT* library. Indeed, we consider the functions:

1. *ot.emd*, which computes the 2-Wasserstein distance with the algorithm from [22], and
2. *ot.bregman.sinkhorn2*, which employs the Sinkhorn-Knopp matrix scaling algorithm from [28] for the entropy regularized optimal transport problem; here, we set the regularization parameter to 0.1 and the stop threshold to 0.001.

In Table 1, the average computational time over ten runs of the three considered numerical schemes is shown; we note that the time is scaled in such a way that the computational time of the approach introduced herein is one unit. As we can clearly observe, the method presented in this work is by magnitudes faster than the two built-in functions from the *POT* library. Rather surprisingly, the *ot.bregman.sinkhorn2* solver took even longer than the *ot.emd* solver and lead to a mean relative error of approximately 0.1956; this is considerably larger than the mean relative error obtained by our deep learning scheme.

We further note that, for this specific experiment, the *total computational time* of training of the neural network from Experiment 6.6 over 100 epochs using the full training set *and* subsequently evaluating the Wasserstein distance function on the test set is *significantly smaller* (the factor being approximately 4.6) than the one of *ot.bregman.sinkhorn2* as specified above, and leads to a smaller error, see Figure 8.

	DNN from Exp. 6.6	ot.emd	ot.bregman.sinkhorn2
Computational time	1	$\approx 1.8767 \cdot 10^4$	$\approx 8.4439 \cdot 10^4$

TABLE 1. Computational time to determine the distance of every element in the test set to the reference element in case of the CIFAR-10 dataset.

We note that in the context of the MNIST and CIFAR-10 datasets, the underlying base sets are discrete. For that reason, we first elaborate analogous results to the ones of Section 3 for this considerably easier case, which however are directly applicable in many applications dealing with a finite dataset.

6.1. Wasserstein distance approximation for a discrete base set

As mentioned before, we consider a Wasserstein Sobolev space over a *finite* metric space \mathfrak{D} . This allows for some simplifications of the results from Section 3, and, in turn, for easier implementation in certain applications; here, we have foremost the grayscale and coloured images from the MNIST and CIFAR-10 datasets in mind. As usual, we signify the set of all probability measures on \mathfrak{D} by

$$\mathcal{P}(\mathfrak{D}) := \left\{ \mu = \sum_{\mathbf{x} \in \mathfrak{D}} i_{\mathbf{x}} \delta_{\mathbf{x}} : \sum_{\mathbf{x} \in \mathfrak{D}} i_{\mathbf{x}} = 1 \text{ and } i_{\mathbf{x}} \in [0, 1] \forall \mathbf{x} \in \mathfrak{D} \right\},$$

where $\delta_{\mathbf{x}}$ denotes the Dirac measure at $\mathbf{x} \in \mathfrak{D}$.

Note that any ordering on the set \mathfrak{D} defines, in the obvious way, a bijection

$$J : \mathfrak{D} \rightarrow \mathbb{S}_{\mathbb{R}^{d_{\mathfrak{D}}}} := \left\{ \mathbf{x} \in \mathbb{R}_{\geq 0}^{d_{\mathfrak{D}}} : \|\mathbf{x}\|_1 = 1 \right\}, \quad (6.1)$$

where $d_{\mathfrak{D}} := |\mathfrak{D}|$ denotes the cardinality of \mathfrak{D} and $\mathbb{R}_{\geq 0} := [0, \infty)$. We show that J is indeed a homeomorphism.

Lemma 6.2. *Given a sequence $(\mu_n)_n \subset \mathcal{P}(\mathfrak{D})$ and an element $\mu \in \mathcal{P}(\mathfrak{D})$, then the following two statements are equivalent:*

- (i) $W_p(\mu_n, \mu) \rightarrow 0$ as $n \rightarrow \infty$;

- (ii) $|J(\mu_n) - J(\mu)| \rightarrow 0$ as $n \rightarrow \infty$, where $|\cdot|$ denotes any p -norm on $\mathbb{R}^{d_{\mathfrak{D}}}$; for instance, we may consider the Euclidean distance $|\cdot| = \|\cdot\|_2$.

In particular, the operator J from (6.1) is a homeomorphism.

Proof. Our proof relies on the fact that $W_p(\mu_n, \mu)$ vanishes as $n \rightarrow \infty$ if and only if μ_n converges weakly to μ , which means that

$$\int_{\mathfrak{D}} \varphi \, d\mu_n \rightarrow \int_{\mathfrak{D}} \varphi \, d\mu \quad \text{for all } \varphi \in C(\mathfrak{D}),$$

where $C(\mathfrak{D})$ denotes the set of (continuous) functions on \mathfrak{D} ; this result can be found, for instance, in [96, Thm. 6.9]. First, let us assume that $\mu_n = \sum_{\mathbf{x} \in \mathfrak{D}} i_{n,\mathbf{x}} \delta_{\mathbf{x}}$ converges to $\mu = \sum_{\mathbf{x} \in \mathfrak{D}} i_{\mathbf{x}} \delta_{\mathbf{x}}$ with respect to the topology induced by the Wasserstein distance W_p , or, equivalently, with respect to the weak topology. For any $\mathbf{x}^* \in \mathfrak{D}$ consider the function $\phi_{\mathbf{x}^*} : \mathfrak{D} \rightarrow \mathbb{R}$ given by

$$\phi_{\mathbf{x}^*}(\mathbf{x}) = \begin{cases} 0 & \mathbf{x} \neq \mathbf{x}^*, \\ 1 & \mathbf{x} = \mathbf{x}^*. \end{cases}$$

Then, thanks to the weak convergence, we have that

$$i_{n,\mathbf{x}^*} = \int_{\mathfrak{D}} \phi_{\mathbf{x}^*} \, d\mu_n \rightarrow \int_{\mathfrak{D}} \phi_{\mathbf{x}^*} \, d\mu = i_{\mathbf{x}^*} \quad \text{as } n \rightarrow \infty.$$

Since this holds for any $\mathbf{x}^* \in \mathfrak{D}$, we indeed have that $|J(\mu_n) - J(\mu)| \rightarrow 0$ as $n \rightarrow \infty$.

Next we assume that (ii) holds true and want to show that this implies (i). So take any $\phi \in C(\mathfrak{D})$, and note that ϕ is uniformly bounded since \mathfrak{D} is finite. Consequently, we find that

$$\left| \int_{\mathfrak{D}} \phi \, d\mu_n - \int_{\mathfrak{D}} \phi \, d\mu \right| \leq \|\phi\|_{\infty} \sum_{\mathbf{x} \in \mathfrak{D}} |i_{n,\mathbf{x}} - i_{\mathbf{x}}| \leq \|\phi\|_{\infty} \sqrt{d_{\mathfrak{D}}} \|J(\mu_n) - J(\mu)\|_2,$$

where we employed the Cauchy–Schwarz inequality in the last step. Since all p -norms on $\mathbb{R}^{d_{\mathfrak{D}}}$ are equivalent, the right-hand side above vanishes as $n \rightarrow \infty$. This proves the weak convergence, and, in turn, the convergence with respect to the W_p -metric. \square

The Heine Borel theorem implies that $\mathbb{S}_{\mathbb{R}^{d_{\mathfrak{D}}}} \subseteq \mathbb{R}^{d_{\mathfrak{D}}}$ is compact, and thus, since $J : \mathcal{P}(\mathfrak{D}) \rightarrow \mathbb{S}_{\mathbb{R}^{d_{\mathfrak{D}}}}$ is a homeomorphism by Lemma 6.2, the same holds true for $\mathcal{P}(\mathfrak{D})$.

Proposition 6.3. *The space $\mathcal{P}(\mathfrak{D})$ endowed with the topology induced by the Wasserstein distance W_p is compact.*

Remark 6.4. As recalled in Section 2.2, the conclusion of Proposition 6.3 is a well-known result and holds true for $\mathcal{P}(K)$, where K is any compact subset of a metric space. This result follows, for instance, from Prokhorov’s theorem. For the sake of completeness, we still wanted to include the proof for the discrete case.

In the following, let $\vartheta \in \mathcal{P}(\mathfrak{D})$ be a reference measure and $\{\mu_k\}_{k=1}^{\infty} \subset \mathcal{P}(\mathfrak{D})$ a countable, dense subset of $\mathcal{P}(\mathfrak{D})$, which exists since $\mathcal{P}(\mathfrak{D})$ is a Polish space. We further denote by (φ_k, ψ_k) the pair of the Kantorovich potentials for ϑ and μ_k ; i.e., we have that, for any $k \in \mathbb{N}$,

$$W_p^p(\vartheta, \mu_k) = \int_{\mathfrak{D}} \varphi_k(\mathbf{x}) \, d\mu_k(\mathbf{x}) + \int_{\mathfrak{D}} \psi_k(\mathbf{x}) \, d\vartheta(\mathbf{x}).$$

It further holds that ψ_k is the c -transform, φ_k^c , of φ_k and vice versa, where

$$\varphi_k^c(\mathbf{x}) := \inf_{\mathbf{y} \in \mathfrak{D}} d(\mathbf{x}, \mathbf{y})^p - \varphi_k(\mathbf{y}), \quad \mathbf{x} \in \mathfrak{D},$$

with $d : \mathfrak{D} \times \mathfrak{D} \rightarrow [0, \infty)$ signifying the metric on \mathfrak{D} ; we refer to [79, Thm. 1.39]. In what follows, for any finite subset $I \subset \mathbb{N}$, let $F_{\vartheta}^{W_p^p}, G_{\vartheta}^I : \mathcal{P}(\mathfrak{D}) \rightarrow \mathbb{R}$ be defined as

$$F_{\vartheta}^{W_p^p}(\mu) := W_p^p(\vartheta, \mu) \quad \text{and} \quad G_{\vartheta}^I(\mu) := \max_{k \in I} \int_{\mathfrak{D}} \varphi_k d\mu + \int_{\mathfrak{D}} \psi_k d\vartheta, \quad (6.2)$$

respectively. The goal is to show that G_{ϑ}^I approximates in a certain sense the Wasserstein distance (function) $F_{\vartheta}^{W_p^p}$.

We first note that $F_{\vartheta}^{W_p^p} : \mathcal{P}(\mathfrak{D}) \rightarrow [0, \infty)$ is continuous: indeed, if $\nu \rightarrow \mu$ in W_p , then

$$\lim_{\nu \rightarrow \mu} F_{\vartheta}^{W_p^p}(\nu) = \lim_{\nu \rightarrow \mu} W_p^p(\nu, \vartheta) = W_p^p(\mu, \vartheta) = F_{\vartheta}^{W_p^p}(\mu).$$

Together with the compactness of the set $\mathcal{P}(\mathfrak{D})$, cf. Proposition 6.3, this immediately implies the following auxiliary result.

Lemma 6.5. *The function $F_{\vartheta}^{W_p^p} : \mathcal{P}(\mathfrak{D}) \rightarrow [0, \infty)$ is uniformly continuous.*

By following along the lines of the proof of [86, Thm. 3.1] one may find that, for any $k \in \mathbb{N}$ and $\mathbf{y}', \mathbf{y}'' \in \mathfrak{D}$,

$$|\varphi_k(\mathbf{y}') - \varphi_k(\mathbf{y}'')| \leq p \max_{\mathbf{x} \in \mathfrak{D}} |d(\mathbf{x}, \mathbf{y}') - d(\mathbf{x}, \mathbf{y}'')| |d(\mathbf{x}, \mathbf{y}')^{p-1} - d(\mathbf{x}, \mathbf{y}'')^{p-1}| \leq Cd(\mathbf{y}', \mathbf{y}''),$$

where the constant C depends on p and \mathfrak{D} , but is independent of k , \mathbf{y}' and \mathbf{y}'' ; for the last inequality we employed the (reverse) triangle inequality and further used that \mathfrak{D} is finite, which, in turn, implies that d is uniformly bounded on $\mathfrak{D} \times \mathfrak{D}$. Then, by similar arguments as in Remark 3.2 and Remark 3.3, we obtain the following bound.

Lemma 6.6. *There exists a positive constant $C_{p, \mathfrak{D}}$, only depending on p , the set \mathfrak{D} , and the base metric d , such that, for any finite subset $I \subset \mathbb{N}$, we have*

$$|G_{\vartheta}^I(\mu) - G_{\vartheta}^I(\nu)| \leq C_{p, \mathfrak{D}} W_p(\mu, \nu) \quad \text{for all } \mu, \nu \in \mathcal{P}(\mathfrak{D}). \quad (6.3)$$

We have gathered all the ingredients to prove the main theorem of the present section.

Theorem 6.7. *For all $\varepsilon > 0$ there exists a finite set $I_{\varepsilon} \subset \mathbb{N}$ such that*

$$\sup_{\mu \in \mathcal{P}(\mathfrak{D})} \left| F_{\vartheta}^{W_p^p}(\mu) - G_{\vartheta}^{I_{\varepsilon}}(\mu) \right| \leq \varepsilon,$$

where $F_{\vartheta}^{W_p^p}$ and $G_{\vartheta}^{I_{\varepsilon}}$ are defined as in (6.2).

Proof. Since $F_{\vartheta}^{W_p^p} : \mathcal{P}(\mathfrak{D}) \rightarrow [0, \infty)$ is uniformly continuous, cf. Lemma 6.5, for all $\varepsilon > 0$ there exists a $\delta_{\varepsilon} > 0$ such that $\left| F_{\vartheta}^{W_p^p}(\mu) - F_{\vartheta}^{W_p^p}(\nu) \right| \leq \varepsilon/2$ whenever $W_p(\mu, \nu) \leq \delta_{\varepsilon}$. Set $\delta := \min \left\{ \delta_{\varepsilon}, \frac{\varepsilon}{2C_{p, \mathfrak{D}}} \right\}$. Then, since $\mathcal{P}(\mathfrak{D})$ is compact by Proposition 6.3, there exists a finite subset $I_{\varepsilon} \subset \mathbb{N}$ such that

$$\mathcal{P}(\mathfrak{D}) = \bigcup_{k \in I_{\varepsilon}} B(\mu_k, \delta), \quad (6.4)$$

where $B(\mu, \delta) = \{\nu \in \mathcal{P}(\mathfrak{D}) : W_p(\mu, \nu) < \delta\}$ is the open ball of radius δ centered at μ . Let $\mu \in \mathcal{P}(\mathfrak{D})$ be arbitrary, and choose $k \in I_{\varepsilon}$ such that $W_p(\mu, \mu_k) < \delta$, cf. (6.4). Then, we have that

$$\begin{aligned} \left| F_{\vartheta}^{W_p^p}(\mu) - G_{\vartheta}^{I_{\varepsilon}}(\mu) \right| &\leq \left| F_{\vartheta}^{W_p^p}(\mu) - F_{\vartheta}^{W_p^p}(\mu_k) \right| + \left| F_{\vartheta}^{W_p^p}(\mu_k) - G_{\vartheta}^{I_{\varepsilon}}(\mu_k) \right| + \left| G_{\vartheta}^{I_{\varepsilon}}(\mu_k) - G_{\vartheta}^{I_{\varepsilon}}(\mu) \right| \\ &\leq \frac{\varepsilon}{2} + \left| F_{\vartheta}^{W_p^p}(\mu_k) - G_{\vartheta}^{I_{\varepsilon}}(\mu_k) \right| + \frac{\varepsilon}{2}, \end{aligned}$$

where the upper bound for the first summand follows from the uniform continuity, and the estimate for the third term is due to (6.3) and the choice of δ . Finally, we note that

$$F_{\vartheta}^{W_p^p}(\mu_k) = W_p^p(\mu_k, \vartheta) = \int_{\mathfrak{D}} \varphi_k d\mu_k + \int_{\mathfrak{D}} \psi_k d\vartheta = G_{\vartheta}^{I_{\varepsilon}}(\mu_k),$$

where the latter equality holds since $k \in I_{\varepsilon}$ and the supremum of the dual problem is attained at the Kantorovich duals (φ_k, ψ_k) . As $\mu \in \mathcal{P}(\mathfrak{D})$ was chosen arbitrarily, the claim is proved. \square

As an immediate consequence of the above theorem we get the ensuing result.

Corollary 6.8. *We have that*

$$\sup_{\mu \in \mathcal{P}(\mathfrak{D})} \left| F_{\vartheta}^{W_p^p}(\mu) - G_{\vartheta}^{1:j}(\mu) \right| \rightarrow 0 \quad \text{as } j \rightarrow \infty,$$

where $1:j := \{1, 2, \dots, j\}$ for $j \in \mathbb{N}$.

6.1.1. *Experiments.* Based on the MNIST and CIFAR-10 datasets, we shall experimentally investigate how well we can approximate the function $F_{\vartheta}^{W_2^2}$ by G_{ϑ} , cf. (6.2), based on the set of pre-computed Kantorovich potentials.

Experiment 6.1. For each

$$j \in \mathcal{J}_{\text{pot}}^M := \{10, 50, 100, 250, 500, 1\,000, 2\,500, 5\,000, 10\,000, 20\,000, 40\,000, 60\,000\}$$

and

$$j \in \mathcal{J}_{\text{pot}}^C := \{10, 50, 100, 250, 500, 1\,000, 2\,500, 5\,000, 10\,000, 20\,000, 30\,000, 40\,000, 50\,000\}$$

for the MNIST and CIFAR-10 datasets, respectively, we randomly picked subsets $I_j \subseteq \{1, 2, \dots, 60\,000\}$ and $I_j \subseteq \{1, 2, \dots, 50\,000\}$, respectively, with $|I_j| = j$, satisfying the hierarchical property $I_j \subseteq I_{j'}$ for $j \leq j'$; here, $|I_j|$ denotes the cardinality of the set I_j . Subsequently, we computed the relative error $\left(F_{\vartheta}^{W_2^2}(\nu) - G_{\vartheta}^{I_j}(\nu) \right) / F_{\vartheta}^{W_2^2}(\nu) > 0$ for one hundred elements $\nu \in \mathcal{T}_{\text{test}}$ drawn from the test set at random, as well as the mean of those individual errors. In Figure 1 we can observe how the individual errors, as well as their mean, decay with an increasing number j of considered potentials. However, in the case of the MNIST dataset (Figure 1 left), even when we consider all the 60 000 potentials related to the entire training set, the mean relative error is still above 0.1. The situation is indeed slightly better for the CIFAR-10 dataset (Figure 1 right), but we still have a mean relative error of approximately 0.02 when all the potentials arising from the entire training set are considered.

Experiment 6.2. We run a similar experiment as the one from before. However, for each $j \in \mathcal{J}_{\text{pot}}^M$ and $j \in \mathcal{J}_{\text{pot}}^C$, respectively, we only consider the mean of the relative errors $\left(F_{\vartheta}^{W_2^2}(\nu) - G_{\vartheta}^{I_j}(\nu) \right) / F_{\vartheta}^{W_2^2}(\nu) > 0$ over all the elements $\nu \in \mathcal{T}_{\text{test}}$. This procedure is repeated twenty times: for each loop the sets I_j are drawn at random once again, but still satisfying the hierarchical property. Subsequently, we plot each of those twenty mean errors, as well as their mean, against the number of pre-computed potentials j ; see Figure 2 (left for MNIST and right for CIFAR-10). For a small number of potentials, the mean error slightly differs for the twenty draws. However, for an increasing number of potentials, they more and more coincide, and finally are equal when all the potentials of the training set are considered. As a comparison, we also depict the mean error obtained when only the potentials corresponding to the digit “1” in the training set for MNIST and the images of ships for CIFAR-10, respectively, are taken into account. In the context of the MNSIT dataset this indeed yields a much larger mean error compared to a comparable number of potentials drawn at random. In the setting of the CIFAR-10 dataset, we only have a relatively small improvement by drawing the potentials over all the possible classes, which is rather surprising.




FIGURE 1. Experiment 6.1: Plot of the relative approximation error $\left(F_{\vartheta}^{W_2^2}(\nu) - G_{\vartheta}^{I_j}(\nu)\right) / F_{\vartheta}^{W_2^2}(\nu)$ for one hundred randomly drawn elements $\nu \in \mathcal{J}_{\text{test}}$ and their mean against the number of potentials j . Left: MNIST. Right: CIFAR-10.

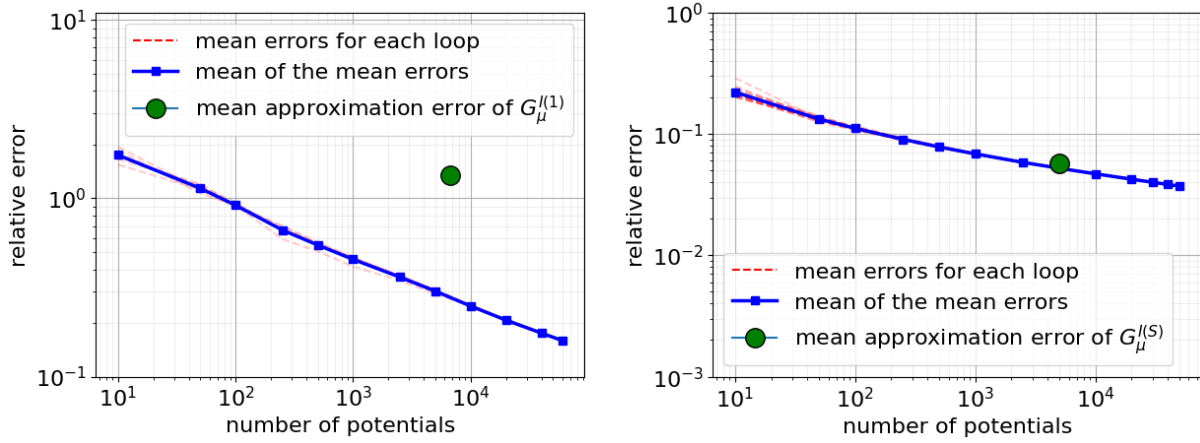


FIGURE 2. Experiment 6.2: Plot of the mean relative errors $\frac{1}{|\mathcal{J}_{\text{test}}|} \sum_{\nu \in \mathcal{J}_{\text{test}}} \left(F_{\vartheta}^{W_2^2}(\nu) - G_{\vartheta}^{I_j}(\nu)\right) / F_{\vartheta}^{W_2^2}(\nu)$ for twenty draws of the index sets I_j and their mean against the number of potentials j . Moreover, we further portray the mean error $\frac{1}{|\mathcal{J}_{\text{test}}|} \sum_{\nu \in \mathcal{J}_{\text{test}}} \left(F_{\vartheta}^{W_2^2}(\nu) - G_{\vartheta}^{I(1)}(\nu)\right) / F_{\vartheta}^{W_2^2}(\nu)$ (left) and $\frac{1}{|\mathcal{J}_{\text{test}}|} \sum_{\nu \in \mathcal{J}_{\text{test}}} \left(F_{\vartheta}^{W_2^2}(\nu) - G_{\vartheta}^{I(S)}(\nu)\right) / F_{\vartheta}^{W_2^2}(\nu)$ (right), where $I(1)$ and $I(S)$ are the sets of indices of the images of the digit “1” and ships, respectively, in the training set of MNIST and CIFAR-10, respectively. Left: MNIST. Right: CIFAR-10.

Experiment 6.3. We repeat Experiment 6.1, but instead of picking one hundred elements from the test set $\mathcal{J}_{\text{test}}$, they are drawn at random from the training set $\mathcal{J}_{\text{train}}$; i.e., we let $\{\mu_{\ell}\}_{\ell \in I_{\text{RT}}} \subset \mathcal{J}_{\text{train}}$, where $I_{\text{RT}} \subset \{1, 2, \dots, 60\,000\}$ for the MNIST dataset and $I_{\text{RT}} \subset \{1, 2, \dots, 50\,000\}$ for the CIFAR-10 dataset, respectively, with $|I_{\text{RT}}| = 100$. Apart from that, and the modified sets

$$\mathcal{J}_{\text{pot}}^{M'} := \{10, 50, 100, 250, 500, 1\,000, 2\,500, 5\,000, 10\,000, 20\,000, 40\,000, 50\,000, 55\,000, 57\,000, 59\,000, 60\,000\},$$

and

$$\mathcal{J}_{\text{pot}}^{C'} := \{10, 50, 100, 250, 500, 1\,000, 2\,500, 5\,000, 10\,000, 20\,000, 30\,000, 40\,000, 45\,000, 47\,000, 49\,000, 50\,000\},$$

respectively, the experimental setup is exactly the same as in Experiment 6.1. We emphasize that if $\ell \in I_j$, then $F_\vartheta^{W_2^2}(\mu_\ell) - G_\vartheta^{I_j}(\mu_\ell) = 0$, and, in turn, $\frac{1}{|I_{\text{RT}}|} \sum_{\ell \in I_{\text{RT}}} \left(F_\vartheta^{W_2^2}(\mu_\ell) - G_\vartheta^{I_{60000}}(\mu_\ell) \right) / F_\vartheta^{W_2^2}(\mu_\ell) = 0$ for the MNIST dataset and $\frac{1}{|I_{\text{RT}}|} \sum_{\ell \in I_{\text{RT}}} \left(F_\vartheta^{W_2^2}(\mu_\ell) - G_\vartheta^{I_{50000}}(\mu_\ell) \right) / F_\vartheta^{W_2^2}(\mu_\ell) = 0$ for the CIFAR-10 dataset, respectively. As we can witness from Figure 3 (left for MNIST and right for CIFAR-10), the individual errors $\left(F_\vartheta^{W_2^2}(\mu_\ell) - G_\vartheta^{I_j}(\mu_\ell) \right) / F_\vartheta^{W_2^2}(\mu_\ell)$ behave qualitatively very similar as for draws from the test set, cp. Figure 1, as long as $\ell \notin I_j$. In particular, the smaller mean error (for a large number of potentials) in the given setting arises due to the elements with zero error contribution.

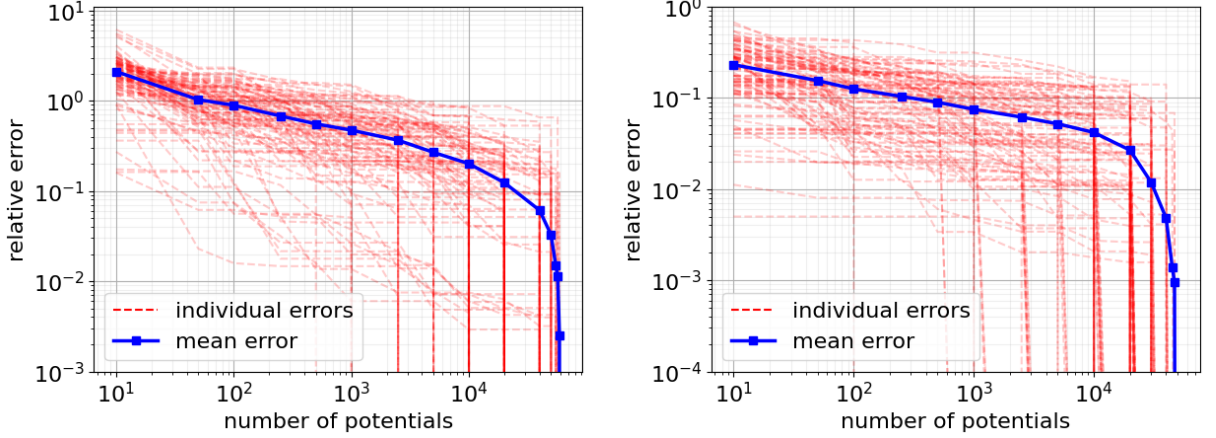


FIGURE 3. Experiment 6.3: Plot of the relative approximation error $\left(F_\vartheta^{W_2^2}(\nu) - G_\vartheta^{I_j}(\nu) \right) / F_\vartheta^{W_2^2}(\nu)$ for one hundred randomly drawn elements $\nu \in \mathcal{T}_{\text{train}}$ and their mean against the number of potentials j . Left: MNIST. Right: CIFAR-10.

Finally, we also repeat the Experiment 6.2, however we replace the test set with 10 000 random elements in the training set $\mathcal{T}_{\text{train}}$. The corresponding error error plot is illustrated in Figure 4.




FIGURE 4. Experiment 6.3 (second part): Based on 10 000 random elements $\{\mu_\ell\}_{\ell \in I_{\text{RT}}} \subset \mathcal{T}_{\text{train}}$, with $|I_{\text{RT}}| = 10000$, we plot the mean relative errors $\frac{1}{|I_{\text{RT}}|} \sum_{\ell \in I_{\text{RT}}} \left(F_\vartheta^{W_2^2}(\mu_\ell) - G_\vartheta^{I_j}(\mu_\ell) \right) / F_\vartheta^{W_2^2}(\mu_\ell)$ for twenty draws of the index sets I_j and their mean against the number of potentials j . Left: MNIST. Right: CIFAR-10.

6.2. Approximation by a deep neural network

In this subsection, we first verify that the function G_ϑ , cf. (6.2), can be realized by a deep neural network, and thereafter highlight that the approximation accuracy from our previous tests can be further enhanced by a suitable training of the network. We first note that, for any pair of continuous functions $\varphi, \psi \in C(\mathfrak{D})$, the mapping

$$\mu \mapsto \int_{\mathfrak{D}} \varphi \, d\mu + \int_{\mathfrak{D}} \psi \, d\vartheta, \quad \mu \in \mathcal{M}(\mathfrak{D}),$$

is affine, where $\mathcal{M}(\mathfrak{D})$ denotes the set of finite, signed measures on \mathfrak{D} . In turn, for any finite subset $I := \{k_1, k_2, \dots, k_j\} \subset \mathbb{N}$, the mapping $L_\vartheta^I : \mathcal{M}(\mathfrak{D}) \rightarrow \mathbb{R}^j$,

$$L_\vartheta^I(\mu) := \left(\int_{\mathfrak{D}} \varphi_{k_1} \, d\mu + \int_{\mathfrak{D}} \psi_{k_1} \, d\vartheta, \dots, \int_{\mathfrak{D}} \varphi_{k_j} \, d\mu + \int_{\mathfrak{D}} \psi_{k_j} \, d\vartheta \right)^\top, \quad (6.5)$$

is affine as well. Moreover, since the extension of the homeomorphism J from (6.1), $\tilde{J} : \mathcal{M}(\mathfrak{D}) \rightarrow \mathbb{R}^{d_\mathfrak{D}}$, is linear, we further have that $\tilde{L}_\vartheta^I := L_\vartheta^I \circ \tilde{J}^{-1} : \mathbb{R}^{d_\mathfrak{D}} \rightarrow \mathbb{R}^j$ is affine, and thus can be represented jointly by a matrix $\mathbf{A}_\vartheta^I \in \mathbb{R}^{j \times d_\mathfrak{D}}$ and a vector $\mathbf{b}_\vartheta^I \in \mathbb{R}^j$. In particular, with the specific numbering of \mathfrak{D} that defines the mapping \tilde{J} , we have that

$$\mathbf{A}_\vartheta^I := \begin{pmatrix} \varphi_{k_1}(\mathbf{x}_1) & \varphi_{k_1}(\mathbf{x}_2) & \cdots & \varphi_{k_1}(\mathbf{x}_{d_\mathfrak{D}}) \\ \varphi_{k_2}(\mathbf{x}_1) & \varphi_{k_2}(\mathbf{x}_2) & & \vdots \\ \vdots & & \ddots & \vdots \\ \varphi_{k_j}(\mathbf{x}_1) & \cdots & \cdots & \varphi_{k_j}(\mathbf{x}_{d_\mathfrak{D}}) \end{pmatrix} \quad \text{and} \quad \mathbf{b}_\vartheta^I := \left(\int_{\mathfrak{D}} \psi_{k_1} \, d\vartheta, \dots, \int_{\mathfrak{D}} \psi_{k_j} \, d\vartheta \right)^\top, \quad (6.6)$$

respectively. We may consider \mathbf{A}_ϑ^I and \mathbf{b}_ϑ^I as the weight matrix and bias vector, respectively, of a shallow neural network with no hidden layers, which represents the mapping \tilde{L}_ϑ^I . We further note that $G_\vartheta^I(\mu) = \max(L_\vartheta^I(\mu)) = \max(\tilde{L}_\vartheta^I(J(\mu)))$, $\mu \in \mathcal{P}(\mathfrak{D})$, where the maximum is taken over the components of the vector $L_\vartheta^I(\mu) = \tilde{L}_\vartheta^I(J(\mu))$; moreover, the composition of functions, which are represented by deep neural networks, is the realization of the concatenation of the corresponding neural networks; we refer, e.g., to [53, Prop. 2.14]. Hence, if we can represent the maxima function

$$(x_1, \dots, x_n)^\top \mapsto \max\{x_1, \dots, x_n\} \quad (6.7)$$

by a deep neural network, then the same holds true for G_ϑ^I . It is indeed known in the literature that the maxima function (6.7) can be realized by a deep ReLU neural network, see, for instance, the works [13, 45, 53].

6.2.1. Deep neural network representation of the maxima function. In the present work, we only recall a (simplified) construction of the maxima function on \mathbb{R}^{2^k} , for some $k \in \mathbb{N}$. For that purpose, we first remind that the ReLU function $\mathfrak{r} : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ is defined by

$$\mathfrak{r}(x) := \max\{0, x\}.$$

In the following, against common convention, we number the weight matrices starting from the output layer towards the input layer; the reason for that is clarified below. Then, the realization of a neural network $\mathcal{A} = (A_l, A_{l-1}, \dots, A_1)$ (without biases) with activation function \mathfrak{r} is given by

$$\mathcal{NN}_{\mathcal{A}}(x) := A_1(\mathfrak{r}(A_2(\mathfrak{r}(\dots(A_{l-1}\mathfrak{r}(A_l x))\dots))),$$

where x is the input of the deep neural network and the activation function \mathfrak{r} is applied componentwise. In the following, we shall construct the neural network representation of the

maxima function (6.7) on \mathbb{R}^n , with $n = 2^k$ for some $k \in \mathbb{N}$. Starting with $k = 1$, the network that realizes (6.7) is given by $\mathcal{A}_1 := (A_2, A_1)$, where

$$A_2 := \begin{pmatrix} 1 & -1 \\ 0 & 1 \\ 0 & -1 \end{pmatrix} \quad \text{and} \quad A_1 := \begin{pmatrix} 1 & 1 & -1 \end{pmatrix}.$$

It is straightforward to check that, for any given $x = (x_1, x_2)^\top \in \mathbb{R}^2$,

$$\mathcal{NN}_{\mathcal{A}_1}(x) = A_1(\mathbf{r}(A_2x)) = \max\{x_1 - x_2, 0\} + \max\{x_2, 0\} - \max\{-x_2, 0\} = \max\{x_1, x_2\}.$$

Now let $B_1 := A_2$ and $C_1 := A_1$. Then, for given $\ell \in \mathbb{N}$, we define the block matrices

$$B_{\ell+1} := \begin{pmatrix} B_\ell & 0 \\ 0 & B_\ell \end{pmatrix} = \underbrace{\begin{pmatrix} A_2 & 0 & \cdots & 0 \\ 0 & A_2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & A_2 \end{pmatrix}}_{2^\ell \text{ times}} \in \mathbb{R}^{3 \cdot 2^\ell \times 2^{\ell+1}}$$

and

$$C_{\ell+1} = \begin{pmatrix} C_\ell & 0 \\ 0 & C_\ell \end{pmatrix} = \underbrace{\begin{pmatrix} A_1 & 0 & \cdots & 0 \\ 0 & A_1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & A_1 \end{pmatrix}}_{2^\ell \text{ times}} \in \mathbb{R}^{2^\ell \times 3 \cdot 2^\ell}.$$

Moreover, still for a natural number $\ell \geq 1$, consider the product

$$D_\ell := B_\ell C_{\ell+1} \in \mathbb{R}^{3 \cdot 2^{\ell-1} \times 3 \cdot 2^\ell}. \quad (6.8)$$

We claim that, for given $k \in \mathbb{N}$, the ReLU neural network $\mathcal{NN}_{\mathcal{A}_k}$ with weight matrices $\mathcal{A}_k := (B_k, D_{k-1}, \dots, D_1, A_1)$ and without biases represents the maxima function on \mathbb{R}^{2^k} , which is shown by induction over k . We already know that the claim holds true for $k = 1$, i.e., for $\mathcal{NN}_{\mathcal{A}_1}$ with $\mathcal{A}_1 = (A_2, A_1)$. So assume that $\mathcal{NN}_{\mathcal{A}_k}$ represents the maxima function for input dimension $n = 2^k$, and show that this implies the claim for $k + 1$. Let $x = (x_1, \dots, x_{2^{k+1}})^\top \in \mathbb{R}^{2^{k+1}}$ be arbitrarily chosen. Then, a simple calculation reveals that

$$\mathcal{NN}_{\mathcal{A}_{k+1}}(x) = \mathcal{NN}_{\mathcal{A}_k} \left((\max\{x_1, x_2\}, \max\{x_3, x_4\}, \dots, \max\{x_{2^{k+1}-1}, x_{2^{k+1}}\})^\top \right),$$

and thus, by the inductive assumption,

$$\begin{aligned} \mathcal{NN}_{\mathcal{A}_{k+1}}(x) &= \max \left\{ \max\{x_1, x_2\}, \max\{x_3, x_4\}, \dots, \max\{x_{2^{k+1}-1}, x_{2^{k+1}}\} \right\} \\ &= \max \{x_1, x_2, \dots, x_{2^{k+1}}\}, \end{aligned}$$

which finishes the inductive argument.

We note that the neural network $\mathcal{NN}_{\mathcal{A}_k}$ is rather sparse and cone shaped, consists of $k = \log_2(n)$ hidden layers, and the i th hidden layer has width $3 \cdot 2^{k-i} = 3n/2^i$ (see Figure 5 for a depiction of the network architecture). The results concerning the architecture of the neural network representing the maxima function for general input dimension $n \in \mathbb{N}$ are very similar; we refer the interested reader to [45, Prop. 5.4].

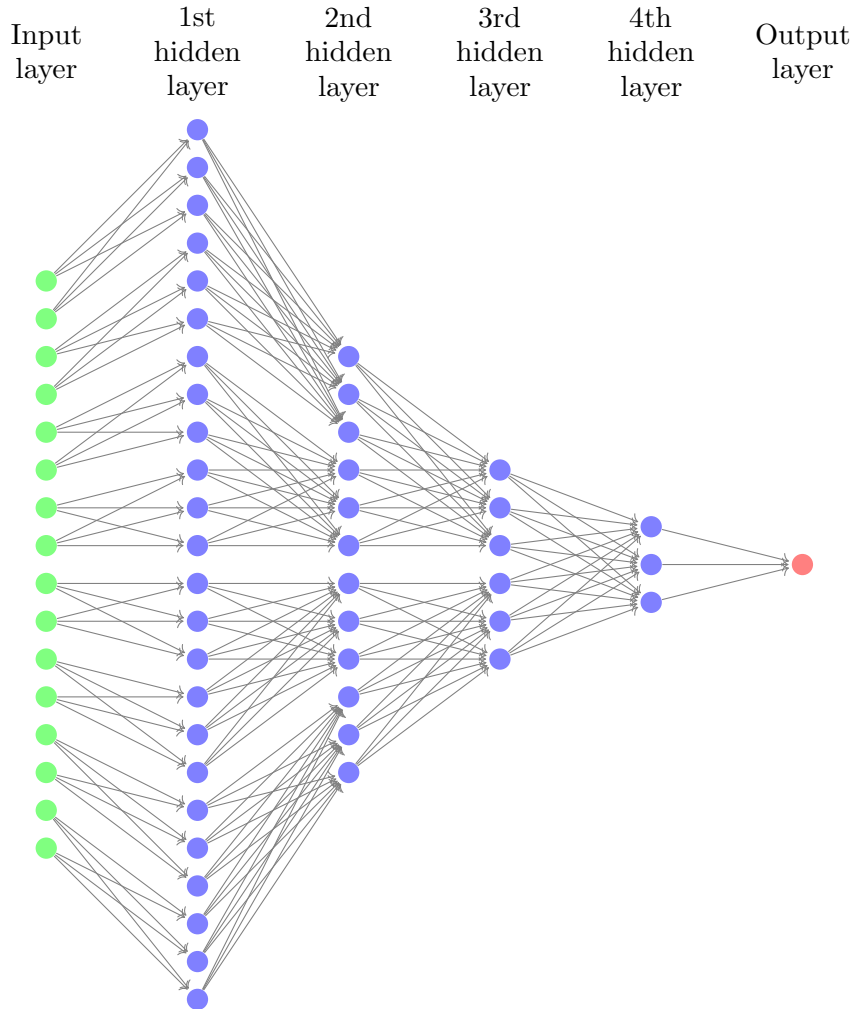


FIGURE 5. The architecture \mathcal{A}_4 , i.e., of the neural network representing the maxima function in \mathbb{R}^{16} .

6.2.2. *Experiments.* We now run some numerical tests to examine the approximation ability of the function $F_{\vartheta}^{W_2^2}$ by our deep neural network architecture introduced above. However, instead of simply computing the weight matrix and bias vector of the first hidden layer, cf. (6.6), with the help of the Python library POT, as (in principle) in the experiments from subsection 6.1, the emphasis in the following tests lies on the training of the deep neural network and the resulting error decay over the epochs. As our loss function we consider the mean absolute error

$$\frac{1}{|\mathcal{B}|} \sum_{\mu \in \mathcal{B}} \left| F_{\vartheta}^{W_2^2}(\mu) - \mathcal{NN}(\mu) \right|, \quad (6.9)$$

where $\mathcal{NN}(\cdot)$ is the realization of the neural network and $\mathcal{B} \subset \mathcal{T}_{\text{train}}$, but still depict the mean relative error

$$\frac{1}{|\mathcal{T}|} \sum_{\mu \in \mathcal{T}} \frac{\left| F_{\vartheta}^{W_2^2}(\mu) - \mathcal{NN}(\mu) \right|}{F_{\vartheta}^{W_2^2}(\mu)}, \quad (6.10)$$

for $\mathcal{T} \in \{\mathcal{T}_{\text{train}}, \mathcal{T}_{\text{test}}\}$, in our figures. For the construction of the neural network we employ the Python open source library PyTorch. Moreover, the network is trained with the ADAM optimizer, see [56] for its original introduction, in combination with mini batches.

Experiment 6.4. In our first experiment, we use a random initialization of the weight matrix and bias vector of the first hidden layer, cf. (6.5) and (6.6); the remaining part of the neural network architecture is constructed in such a way that it represents the maxima function, see Section 6.2.1, with fixed weights and biases. Here, we consider 2^{12} neurons in the first hidden layer for both the MNIST and CIFAR-10 datasets, resulting in 3 211 264 and 12 587 008 trainable parameters, respectively. As we can see in Figure 6 (left), for the MNIST dataset we have a step decay of the mean relative error for both the training set as well as the test set at the beginning of the training. Yet, after this short initial phase of substantial drop in the loss, we can only observe a slight improvement thereafter. However, at this point, the error is already well below $0.1594\dots$, which is the mean relative error obtained in Experiment 6.2 when all the potentials corresponding to the training set $\mathcal{T}_{\text{train}}$ are considered. This is in contrast to the corresponding experiment for the CIFAR-10 dataset, see Figure 6 (right). Here, the error only improves little over the three hundred training epochs, and even the final mean relative error is remarkably larger than the error in Experiment 6.2 for 50 000 pre-computed potentials arising from the training set, cf. Figure 2 (right).




FIGURE 6. Experiment 6.4: Approximation of the function $F_{\vartheta}^{W_2^2}$ by a deep neural network with 2^{12} hidden neurons in the first hidden layer, corresponding to the trainable part of the deep neural network. The red dotted line indicates the final error when all the potentials from the training set were considered in Experiment 6.2. Left: MNIST. Right: CIFAR-10.

Experiment 6.5. As can be observed in the previous experiment, the deep neural network approach performs worse than the one based on pre-computed potentials in the case of the CIFAR-10 dataset. In this experiment, we shall investigate if a combination of those two ideas may further enhance the approximation accuracy. In particular, we initialize the weight matrix and bias vector, corresponding to the mapping from the input layer to the first hidden layer, with 2^{12} pre-computed potentials; cf. (6.5), (6.6). The remaining part of the architecture represents the maxima function, and, as before, has fixed weights and biases. Subsequently, we train the deep neural network over one hundred epochs, using ADAM combined with mini batches. As we can observe from Figure 7, the mean error for both the training set as well as the test set drastically drops at the initial phase of the training. Thereafter, for both the MNIST and

CIFAR-10 datasets, the training loss still steadily decays, however, with a reduced rate. Albeit, for CIFAR-10, the test loss stagnates after the initial phase, see Figure 7 (right).




FIGURE 7. Experiment 6.5. Approximation of the function $F_{\vartheta}^{W_2^2}$ by a deep neural network with 2^{12} hidden neurons in the first hidden layer, with the weight matrix and bias vector corresponding to (6.5) and (6.6) initialized by pre-computed Kantorovich potentials. The red dotted line indicates the final error when all the potentials from the training set were considered in Experiment 6.2. Left: MNIST. Right: CIFAR-10.

Experiment 6.6. We use the same neural network design as in the previous experiments; i.e., we consider a weight matrix and bias vector corresponding to (6.5), (6.6) between the input and first hidden layers (without ReLU activation function), and the remainder of the architecture represents the maxima function as outlined in Section 6.2.1. However, in contrast to the experiments before, we allow all the neural network parameters to be trainable, including the ones from the part corresponding to the representation of the maxima function. Moreover, in one run the entries of the weight matrix and bias vector corresponding to (6.6) are randomly initialized according to the default setting of the PyTorch package (see Figure 8 (left)), and the other time we use pre-computed Kantorovich potentials as in Experiment 6.5; here, we only consider the CIFAR-10 dataset, and the first hidden layer consists of 2^{10} neurons. Figure 8 indicates that this might further enhance the accuracy compared to the setting of Experiment 6.5, mostly for the generalization; we emphasize that in the present numerical test the number of trainable parameters is 6 294 010 compared to 12 587 008 in the previous experiment. Surprisingly, the performance is even slightly superior for the random initialization of the weight matrix and the bias vector of the first hidden layer compared to the initialization with pre-computed potentials.

Experiment 6.7. We reconsider the fully trainable neural network from Experiment 6.6 with random initialization of the matrix and bias vector of the first hidden layer. However, we regularise the loss function (6.9) with the discretised $H^{1,2}$ -norm from (5.14); i.e., the loss function takes the form

$$\frac{1}{|\mathcal{B}|} \left(\sum_{\mu \in \mathcal{B}} |F_{\vartheta}^{W_2^2}(\mu) - \text{NN}(\mu)| + \lambda \left(|\text{NN}(\mu)|^2 + \int_{\mathbb{R}^d} |\text{DNN}(x, \mu)|^2 d\mu(x) \right) \right). \quad (6.11)$$

We recall that the neural network resembles a cylinder function so that DNN is well-defined, cf. (2.11), and can be approximated by means of a finite difference scheme for the space derivatives.




FIGURE 8. Experiment 6.6: Approximation of the function $F_{\vartheta}^{W_2^2}$ by a fully trainable deep neural network with 2^{10} hidden neurons in the first hidden layer, composed with the neural network that initially, i.e., before the training, represents the maxima function. The red dotted line indicates the final error when all the potentials from the training set were considered in Experiment 6.2. Left: Random initialization of the weight matrix and bias vector of the first hidden layer. Right: Initialization based on pre-computed potentials.

In Figure 9 we plot the evolution of the mean relative error (6.10), for $\mathcal{T}_{\text{train}}$ and $\mathcal{T}_{\text{test}}$, respectively, over the training epochs. For both choices of the regularization parameter, $\lambda = 0.1$ (left) and $\lambda = 0.001$ (right), respectively, the mean relative error behaves similarly as in the case of the unregularized loss function, cf. Figure 8 (left).




FIGURE 9. Experiment 6.7: Repetition of the Experiment 6.6, however with an additional regularization term in the loss function, see (6.11). Left: $\lambda = 0.1$. Right: $\lambda = 0.001$.

6.3. Adversarial deep neural network approach for the solution of the saddle-point problem

In the following, we show how the saddle-point problem (5.13) can be leveraged to approximate the minimizer of (5.1) in applications. As in Remark 5.4, let

$$\mathfrak{p}_N = \frac{1}{N} \sum_{i=1}^N \delta_{\mu_i}$$

be our random empirical measure. As already pointed out in that remark, it is not always easy to find a best approximating finite dimensional subspace \mathcal{V}_n of cylinder functions. However, we know that cylinder functions can be modelled by deep neural networks; cf. Section 6.2. So let F_θ be a cylinder function that is realized by a deep neural network, where θ denotes the specific network parameters stemming from an admissible set Θ of trainable parameters; this serves as our solution network. In particular, the infimum in (5.13) is taken over

$$\{F_\theta : \theta \in \Theta\} \subset \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d)).$$

Similarly we consider an adversarial network H_ξ that realizes a cylinder function, and which depends on the trainable parameters $\xi \in \Xi$, for the supremum in (5.13). This gives rise to the saddle point problem

$$\inf_{\theta \in \Theta} \sup_{\xi \in \Xi} \frac{(F_\theta - \tilde{F}, H_\xi)_{L_{\mathfrak{p}_N}^2} + \lambda \text{pCE}_{2, \mathfrak{p}_N}(F_\theta, H_\xi)}{|H_\xi|_{pH_{\mathfrak{p}_N}^{1,2}}}. \quad (6.12)$$

We shall mention that a methodologically equivalent approach was introduced in [99] for the numerical solution of the weak formulation of high-dimensional partial differential equations, which often arise as the Euler–Lagrange equation of an energy minimization problem. Moreover, this is closely related to the main idea of *generative adversarial networks* as introduced in [44].

Remark 6.9. Even if we consider the same architecture for the solution and adversarial networks, respectively, we may not find a parameter $\theta \in \Theta$ in the admissible set such that

$$(F_\theta - \tilde{F}, H_\xi)_{L_{\mathfrak{p}_N}^2} + \lambda \text{pCE}_{2, \mathfrak{p}_N}(F_\theta, H_\xi) = 0 \quad \text{for all } \xi \in \Xi;$$

cp. (5.7). This is due to the fact that

$$\{F_\theta : \theta \in \Theta\} \subset \mathfrak{C}(\mathcal{P}(\mathbb{R}^d), C_b^1(\mathbb{R}^d))$$

is not a linear subspace.

We now discuss the specific algorithm, cf. Algorithm 1, which is very closely related to the one from [99], to obtain a suitable solution approximation. For given neural network architectures for F_θ and H_ξ , we initialize the weights $\theta \in \Theta$ and $\xi \in \Xi$. Subsequently, we repeat the following procedure until a stopping criterion, which could be a prescribed number of loops or an error tolerance for the loss function $L_\xi(\theta)$, cf. (6.14) below, is satisfied:

(1) for a fixed parameter $\theta \in \Theta$, we define the adversarial loss function

$$L_\theta(\xi) := - \frac{(F_\theta - \tilde{F}, H_\xi)_{L_{\mathfrak{p}_N}^2} + \lambda \text{pCE}_{2, \mathfrak{p}_N}(F_\theta, H_\xi)}{|H_\xi|_{pH_{\mathfrak{p}_N}^{1,2}}} \quad (6.13)$$

with input variable $\xi \in \Xi$. Then we train, for instance by applying gradient descent, the neural network H_ξ over a prescribed number of epochs $N_\Xi \geq 1$ to minimize the loss function $L_\theta(\xi)$;

Algorithm 1 Adversarial deep neural network algorithm for the solution of the Euler–Lagrange equation

Input: Training data $\{(\mu_j, y_j)\}_{j=1}^N$ and number of epochs $N_{\Xi}, N_{\Theta} \geq 1$.

Initialize: Neural networks F_{θ} and H_{ξ} .

```

1: repeat
2:   for  $k = 1, \dots, N_{\Xi}$  do
3:     One training step for  $H_{\xi}$  to minimize the loss function  $L_{\theta}(\xi)$  from (6.13).
4:   end for
5:   for  $k = 1, \dots, N_{\Theta}$  do
6:     One training step for  $F_{\theta}$  to minimize the loss function  $L_{\xi}(\theta)$  from (6.14).
7:   end for
8: until stopping criterion is satisfied

```

Output: Trained solution network F_{θ} .

(2) for given $\xi \in \Xi$, we define the solution loss function

$$L_{\xi}(\theta) := \frac{(F_{\theta} - \tilde{F}, H_{\xi})_{L^2_{\mathfrak{p}_N}} + \lambda \text{pCE}_{2, \mathfrak{p}_N}(F_{\theta}, H_{\xi})}{|H_{\xi}|_{\mathcal{H}^{1,2}_{\mathfrak{p}_N}}}, \quad (6.14)$$

which depends on $\theta \in \Theta$. Similar as in (1), we employ a training procedure for the neural network F_{θ} to optimize the loss function (6.14).

Remark 6.10. We note that both loss functions L_{θ} and L_{ξ} , cf. (6.14) and (6.13), respectively, depend on a set of given data $\{(\mu_j, y_j)\}_{j=1}^N$, which is our training set, through the empirical measure \mathfrak{p}_N and the corresponding noisy function values. In practice, we rather use mini batches thereof for the training, and apply optimization procedures that are superior to gradient descent, such as stochastic gradient descent or ADAM (see [56]).

6.3.1. *Experiments.* We now test our Algorithm 1 in the context of the MNIST dataset. We once more use the network architecture from Experiment 6.6, adapted to the smaller input dimension of the MNIST dataset compared to CIFAR-10, for both the solution network as well as the adversarial network.

Experiment 6.8. We compare the performance of Algorithm 1, with the regularization parameter set to $\lambda = 0.001$, for two distinct choices of the pairs of parameters (N_{Ξ}, N_{Θ}) . In Figure 10 (left) we set $N_{\Xi} = 1$ and $N_{\Theta} = 2$, whereas in Figure 10 (right) we choose $N_{\Xi} = 2$ and $N_{\Theta} = 1$. We can observe in both figures that the mean relative error of both the training as well as the test set do nicely decay, with a comparable overall performance as in Experiment 6.5. Even though we obtain, in general, a better approximation accuracy with the approach outlined in Section 6.2, it is noteworthy that the mean relative error of the training and test sets almost coincide in this experiment, in contrast to the ones in Section 6.2, meaning that the generalization error almost vanishes.

Experiment 6.9. We again consider the Euler–Lagrange approach from Section 5, however we work on $L^2(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathfrak{m})$ rather than $H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathfrak{m})$. In turn, the discrete saddle point problem, cf. (6.12), reads as

$$\inf_{\theta \in \Theta} \sup_{\xi \in \Xi} \frac{(F_{\theta} - \tilde{F}, H_{\xi})_{L^2_{\mathfrak{p}_N}}}{|H_{\xi}|_{L^2_{\mathfrak{p}_N}}},$$

where

$$|H|_{L^2_{\mathfrak{p}_N}}^2 := (H, H)_{L^2_{\mathfrak{p}_N}} = \frac{1}{N} \sum_{j=1}^N |H(\mu_j)|^2, \quad H \in L^2(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathfrak{m}).$$




FIGURE 10. Experiment 6.8: Decay of the mean relative error for the MNIST dataset based on Algorithm 1. Left: $N_{\Xi} = 1$ and $N_{\Theta} = 2$. Right: $N_{\Xi} = 2$ and $N_{\Theta} = 1$.

Except for that, we use the same framework as in Experiment 6.8. In Figure 11 we can once more observe the decay of the mean relative error, with a similar overall behaviour of the error over the epochs as before. Albeit, the error seems to fluctuate more than in Experiment 6.8; i.e., considering the $H^{1,2}(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathfrak{m})$ framework instead of working on $L^2(\mathcal{P}_2(\mathbb{R}^d), W_2, \mathfrak{m})$ might indeed lead to a smoother approximation of the target function.




FIGURE 11. Experiment 6.9: Decay of the mean relative error for the MNIST dataset based on Algorithm 1. Left: $N_{\Xi} = 1$ and $N_{\Theta} = 2$. Right: $N_{\Xi} = 2$ and $N_{\Theta} = 1$.

REFERENCES

- [1] Luigi Ambrosio, Elia Brué, and Daniele Semola. *Lectures on optimal transport*, volume 130 of *Unitext*. Springer, Cham, 2021.
- [2] Luigi Ambrosio, Matthias Erbar, and Giuseppe Savaré. Optimal transport, Cheeger energies and contractivity of dynamic transport distances in extended spaces. *Nonlinear Anal.*, 137:77–134, 2016.
- [3] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows in metric spaces and in the space of probability measures*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, second edition, 2008.
- [4] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. Calculus and heat flow in metric measure spaces and applications to spaces with Ricci bounds from below. *Invent. Math.*, 195(2):289–391, 2014.

- [5] Luigi Ambrosio, Shouhei Honda, Jacobus W. Portegies, and David Tewodrose. Embedding of $RCD(K,N)$ spaces in L^2 via eigenfunctions. *Journal of Functional Analysis*, 280(10):108968, 2021.
- [6] Sigurd Angenent, Steven Haker, and Allen Tannenbaum. Minimizing flows for the Monge-Kantorovich problem. *SIAM J. Math. Anal.*, 35(1):61–97, 2003.
- [7] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223. PMLR, 06–11 Aug 2017.
- [8] H. Attouch. *Variational convergence for functions and operators*. Applicable Mathematics Series. Pitman (Advanced Publishing Program), Boston, MA, 1984.
- [9] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.*, 16(1):78–96, 1987.
- [10] Mohit Bansil and Jun Kitagawa. A newton algorithm for semidiscrete optimal transport with storage fees. *SIAM Journal on Optimization*, 31(4):2586–2613, 2021.
- [11] Federico Basseti, Stefano Gualandi, and Marco Veneroni. On the computation of Kantorovich-Wasserstein distances between two-dimensional histograms by uncapacitated minimum cost flows. *SIAM J. Optim.*, 30(3):2441–2469, 2020.
- [12] Saurav Basu, Soheil Kolouri, and Gustavo K. Rohde. Detecting and visualizing cell phenotype differences from microscopy images using transport-based morphometry. *Proceedings of the National Academy of Sciences*, 111(9):3448–3453, 2014.
- [13] Christian Beck, Arnulf Jentzen, and Benno Kuckuck. Full error analysis for the training of deep neural networks. *Infin. Dimens. Anal. Quantum Probab. Relat. Top.*, 25(2):Paper No. 2150020, 76, 2022.
- [14] Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numer. Math.*, 84(3):375–393, 2000.
- [15] Jean-David Benamou, Yann Brenier, and Kevin Guittet. Numerical analysis of a multi-phasic mass transport problem. In *Recent advances in the theory and applications of mass transport*, volume 353 of *Contemp. Math.*, pages 1–17. Amer. Math. Soc., Providence, RI, 2004.
- [16] Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. Iterative bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.
- [17] Dimitri P. Bertsekas. A new algorithm for the assignment problem. *Math. Programming*, 21(2):152–171, 1981.
- [18] Dimitri P. Bertsekas and David A. Castañón. The auction algorithm for the transportation problem. volume 20, pages 67–96. 1989. Network optimization and applications.
- [19] Dimitri P. Bertsekas and Jonathan Eckstein. Dual coordinate step methods for linear network flow problems. *Math. Programming*, 42(2, (Ser. B)):203–243, 1988.
- [20] Anders Björn and Jana Björn. *Nonlinear potential theory on metric spaces*, volume 17 of *EMS Tracts in Mathematics*. European Mathematical Society (EMS), Zürich, 2011.
- [21] Helmut Bölcskei, Philipp Grohs, Gitta Kutyniok, and Philipp Petersen. Optimal approximation with sparsely connected deep neural networks. *SIAM Journal on Mathematics of Data Science*, 1(1):8–45, 2019.
- [22] Nicolas Bonneel, Michiel Panne, Sylvain Paris, and Wolfgang Heidrich. Displacement interpolation using lagrangian mass transport. *ACM Trans. Graph.*, 30:158, 12 2011.
- [23] Rick Chartrand, Brendt Wohlberg, Kevin Vixie, and Erik Bollt. A gradient descent solution to the monge-kantorovich problem. *Applied Mathematical Sciences*, 3:1071–1080, 01 2009.
- [24] Jeff Cheeger. Differentiability of Lipschitz functions on metric measure spaces. *Geom. Funct. Anal.*, 9(3):428–517, 1999.
- [25] Lénaïc Chizat, Pierre Roussillon, Flavien Léger, François-Xavier Vialard, and Gabriel Peyré. Faster wasserstein distance estimation with the sinkhorn divergence. NIPS’20, Red Hook, NY, USA, 2020. Curran Associates Inc.
- [26] Albert Cohen, Mark A. Davenport, and Dany Leviatan. Correction to: On the stability and accuracy of least squares approximations [MR3105946]. *Found. Comput. Math.*, 19(1):239, 2019.
- [27] Nicolas Courty, Rémi Flamary, and Mélanie Ducoffe. Learning wasserstein embeddings. In *International Conference on Learning Representations*, 2018.
- [28] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- [29] George V. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2:303–314, 1989.
- [30] Gianni Dal Maso. *An introduction to Γ -convergence*, volume 8 of *Progress in Nonlinear Differential Equations and their Applications*. Birkhäuser Boston, Inc., Boston, MA, 1993.

- [31] Lorenzo Dello Schiavo. A Rademacher-type theorem on L^2 -Wasserstein spaces over closed Riemannian manifolds. *J. Funct. Anal.*, 278(6):108397, 57, 2020.
- [32] W. Edwards Deming and Frederick F. Stephan. On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *Ann. Math. Statistics*, 11:427–444, 1940.
- [33] Ronald DeVore, Boris Hanin, and Guergana Petrova. Neural network approximation. *Acta Numerica*, 30:327–444, 2021.
- [34] Dennis Elbrächter, Dmytro Perekrestenko, Philipp Grohs, and Helmut Bölcskei. Deep neural network approximation theory. *IEEE Transactions on Information Theory*, 67:2581–2623, 2019.
- [35] Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z. Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, Léo Gautheron, Nathalie T.H. Gayraud, Hicham Janati, Alain Rakotomamonjy, Ievgen Redko, Antoine Rolet, Antony Schutz, Vivien Seguy, Danica J. Sutherland, Romain Tavenard, Alexander Tong, and Titouan Vayer. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.
- [36] Massimo Fornasier, Pascal Heid, and Giacomo E. Sodini. Stability and accuracy of Tikhonov regularization on infinitesimally Hilbertian metric measure spaces. in preparation.
- [37] Massimo Fornasier, Giacomo E. Sodini, and Giuseppe Savaré. Density of subalgebras of Lipschitz functions in metric Sobolev spaces and applications to Sobolev-Wasserstein spaces. *arXiv:2209.00974*, 2022.
- [38] Joel Franklin and Jens Lorenz. On the scaling of multidimensional matrices. *Linear Algebra Appl.*, 114/115:717–735, 1989.
- [39] Aude Genevay, Marco Cuturi, Gabriel Peyré, and Francis Bach. Stochastic optimization for large-scale optimal transport. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS’16, page 3440–3448, Red Hook, NY, USA, 2016. Curran Associates Inc.
- [40] Nicola Gigli. On the differential structure of metric measure spaces and applications. *Mem. Amer. Math. Soc.*, 236(1113):vi+91, 2015.
- [41] Nicola Gigli, Andrea Mondino, and Tapio Rajala. Euclidean spaces as weak tangents of infinitesimally Hilbertian metric measure spaces with Ricci curvature bounded below. *J. Reine Angew. Math.*, 705:233–244, 2015.
- [42] Nicola Gigli and Enrico Pasqualetto. *Lectures on nonsmooth differential geometry*, volume 2 of *SISSA Springer Series*. Springer, Cham, [2020] ©2020.
- [43] Andrew V. Goldberg and Robert E. Tarjan. Finding minimum-cost circulations by canceling negative cycles. *J. Assoc. Comput. Mach.*, 36(4):873–886, 1989.
- [44] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [45] Philipp Grohs, Arnulf Jentzen, and Diyora Salimova. Deep neural network approximations for solutions of PDEs based on Monte Carlo algorithms. *Partial Differ. Equ. Appl.*, 3(4):Paper No. 45, 41, 2022.
- [46] K. Guittet. On the time-continuous mass transport problem and its approximation by augmented Lagrangian techniques. *SIAM J. Numer. Anal.*, 41(1):382–399, 2003.
- [47] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [48] Steven Haker, Lei Zhu, Allen Tannenbaum, and Sigurd Angenent. Optimal mass transport for registration and warping. *Int. J. Comput. Vision*, 60(3):225–240, dec 2004.
- [49] Lukas Herrmann, Joost A. A. Opschoor, and Christoph Schwab. Constructive deep ReLU neural network approximation. *J. Sci. Comput.*, 90(2):Paper No. 75, 37, 2022.
- [50] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257, 1991.
- [51] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.
- [52] Zhangkai Huang. Isometric immersions of $\text{rcd}(k, n)$ spaces via heat kernels. *Calculus of Variations and Partial Differential Equations*, 62(4):121, 2023.
- [53] Arnulf Jentzen and Adrian Riekert. Strong overall error analysis for the training of artificial neural networks via random initializations. *arXiv:2012.08443*, 2020.
- [54] Keisuke Kawano, Satoshi Koide, and Takuro Kutsuna. Learning wasserstein isometric embedding for point clouds. In *2020 International Conference on 3D Vision (3DV)*, pages 473–482, 2020.

- [55] Patrick Kidger and Terry Lyons. Universal Approximation with Deep Narrow Networks. In Jacob Abernethy and Shivani Agarwal, editors, *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 2306–2327. PMLR, 09–12 Jul 2020.
- [56] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- [57] Jun Kitagawa, Quentin Mérigot, and Boris Thibert. A newton algorithm for semi-discrete optimal transport. *Journal of the European Mathematical Society*, 21, 03 2016.
- [58] Jun Kitagawa, Quentin Mérigot, and Boris Thibert. Convergence of a Newton algorithm for semi-discrete optimal transport. *J. Eur. Math. Soc.*, 21(9):2603–2651, 2019.
- [59] A. N. Kolmogorov. On the representation of continuous functions of many variables by superposition of continuous functions of one variable and addition. *Dokl. Akad. Nauk SSSR*, 114:953–956, 1957.
- [60] Soheil Kolouri, Navid Naderialzadeh, Gustavo K. Rohde, and Heiko Hoffmann. Wasserstein embedding for graph learning. In *International Conference on Learning Representations*, 2021.
- [61] Soheil Kolouri, Se Rim Park, Matthew Thorpe, Dejan Slepcev, and Gustavo K. Rohde. Optimal mass transport: Signal processing and machine-learning applications. *IEEE Signal Processing Magazine*, 34(4):43–59, 2017.
- [62] Soheil Kolouri, Akif B. Tosun, John A. Ozolek, and Gustavo K. Rohde. A continuous linear optimal transport approach for pattern analysis in image datasets. *Pattern Recognition*, 51:453–462, 2016.
- [63] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- [64] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Res. Logist. Quart.*, 2:83–97, 1955.
- [65] Bruno Lévy. A numerical algorithm for L_2 semi-discrete optimal transport in 3D. *ESAIM Math. Model. Numer. Anal.*, 49(6):1693–1715, 2015.
- [66] Matthias Liero, Alexander Mielke, and Giuseppe Savaré. Optimal entropy-transport problems and a new Hellinger-Kantorovich distance between positive measures. *Invent. Math.*, 211(3):969–1117, 2018.
- [67] Haibin Ling and Kazunori Okada. An efficient earth mover’s distance algorithm for robust histogram comparison. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29:840–853, 2007.
- [68] Huidong Liu, Xianfeng GU, and Dimitris Samaras. A two-step computation of the exact GAN Wasserstein distance. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3159–3168. PMLR, 10–15 Jul 2018.
- [69] Grégoire Loeper and Francesca Rapetti. Numerical solution of the Monge-Ampère equation by a Newton’s algorithm. *C. R. Math. Acad. Sci. Paris*, 340(4):319–324, 2005.
- [70] John Lott and Cédric Villani. Ricci curvature for metric-measure spaces via optimal transport. *ArXiv Mathematics e-prints*, December 2004.
- [71] Danka Lučić and Enrico Pasqualetto. Gamma-convergence of Cheeger energies with respect to increasing distances. *J. Math. Anal. Appl.*, 515(1):Paper No. 126415, 10, 2022.
- [72] Bruno Lévy and Erica L. Schwindt. Notions of optimal transport theory and how to implement them on a computer. *Computers & Graphics*, 72:135–148, 2018.
- [73] Anton Mallasto, Guido Montúfar, and Augusto Gerolin. How well do wgens estimate the wasserstein metric? *arXiv:1910.03875*, 2019.
- [74] Quentin Mérigot. A multiscale approach to optimal transport. *Computer Graphics Forum*, 30(5):1583–1592, 2011.
- [75] Ofir Pele and Michael Werman. Fast and robust earth mover’s distances, 2009.
- [76] Gabriel Peyré and Marco Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [77] Aaditya Ramdas, Nicolás García Trillos, and Marco Cuturi. On wasserstein two-sample testing and related families of nonparametric tests. *Entropy*, 19(2):47, 2017.
- [78] Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision*, 40:99–121, 2000.
- [79] Filippo Santambrogio. *Optimal transport for applied mathematicians*, volume 87 of *Progress in Nonlinear Differential Equations and their Applications*. Birkhäuser/Springer, Cham, 2015. Calculus of variations, PDEs, and modeling.
- [80] Giuseppe Savaré. Sobolev spaces in extended metric-measure spaces. In *New Trends on Analysis and Geometry in Metric Spaces*, pages 117–276. Springer, 2022.
- [81] Bernhard Schmitzer. Stabilized sparse scaling algorithms for entropy regularized transport problems. *SIAM Journal on Scientific Computing*, 41(3):A1443–A1481, 2019.
- [82] Jörn Schrieber. *Algorithms for Optimal Transport and Wasserstein Distances*. PhD thesis, Georg-August-Universität Göttingen, 2019.

- [83] Vivien Seguy and Marco Cuturi. Principal geodesic analysis for probability measures under the optimal transport metric. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.
- [84] Nageswari Shanmugalingam. Newtonian spaces: an extension of Sobolev spaces to metric measure spaces. *Rev. Mat. Iberoamericana*, 16(2):243–279, 2000.
- [85] Richard Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *Ann. Math. Statist.*, 35:876–879, 1964.
- [86] Giacomo E. Sodini. The general class of Wasserstein Sobolev spaces: density of cylinder functions, reflexivity, uniform convexity and Clarkson’s inequalities. *arXiv:2212.10955*, 2022.
- [87] Giacomo Enrico Sodini and Lorenzo Dello Schiavo. Universal infinitesimal hilbertianity of the space of measures with the hellinger-kantorovich distance and applications. *In preparation*, 2024.
- [88] Mathias Sonleitner and Mario Ullrich. On the power of iid information for linear approximation. *arXiv:2310.12740*, 2023.
- [89] Jan Stanczuk, Christian Etmann, Lisa Maria Kreusser, and Carola-Bibiane Schönlieb. Wasserstein gans work because they fail (to approximate the wasserstein distance). *arXiv:2103.01678*, 2021.
- [90] Karl-Theodor Sturm. On the geometry of metric measure spaces. I. *Acta Math.*, 196(1):65–131, 2006.
- [91] Karl-Theodor Sturm. On the geometry of metric measure spaces. II. *Acta Math.*, 196(1):133–177, 2006.
- [92] Robert E. Tarjan. Dynamic trees as search trees via Euler tours, applied to the network simplex algorithm. volume 78, pages 169–177. 1997. *Network optimization: algorithms and applications* (San Miniato, 1993).
- [93] David Tewodrose. A survey on spectral embeddings and their application in data analysis. *Séminaire de théorie spectrale et géométrie*, 35:197–244, 2017-2019.
- [94] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.*, 12(4):389–434, 2012.
- [95] Cédric Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2003.
- [96] Cédric Villani. *Optimal transport*, volume 338 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2009. Old and new.
- [97] Cédric Villani. *Optimal transport. Old and new*, volume 338 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 2009.
- [98] Wei Wang, Dejan Slepčev, Saurav Basu, John Ozolek, and Gustavo Rohde. A linear optimal transportation framework for quantifying and visualizing variations in sets of images. *International journal of computer vision*, 101:254–269, 06 2013.
- [99] Yaohua Zang, Gang Bao, Xiaojing Ye, and Haomin Zhou. Weak adversarial networks for high-dimensional partial differential equations. *Journal of Computational Physics*, 411:109409, 2020.
- [100] E. Zeidler. *Nonlinear functional analysis and its applications. II/B*. Springer-Verlag, New York, 1990.

MASSIMO FORNASIER: TUM SCHOOL OF COMPUTATION, INFORMATION, AND TECHNOLOGY - DEPARTMENT OF MATHEMATICS, BOLTZMANNSTRASSE 3, 85748 GARCHING BEI MÜNCHEN (GERMANY)
 & TUM INSTITUTE FOR ADVANCED STUDIES
 & MUNICH DATA SCIENCE INSTITUTE
 & MUNICH CENTER FOR MACHINE LEARNING
Email address: `massimo.fornasier@cit.tum.de`

PASCAL HEID: TUM SCHOOL OF COMPUTATION, INFORMATION, AND TECHNOLOGY - DEPARTMENT OF MATHEMATICS, BOLTZMANNSTRASSE 3, 85748 GARCHING BEI MÜNCHEN (GERMANY)
 & MUNICH CENTER FOR MACHINE LEARNING
Email address: `hepa@ma.tum.de`

GIACOMO ENRICO SODINI: INSTITUT FÜR MATHEMATIK - FAKULTÄT FÜR MATHEMATIK - UNIVERSITÄT WIEN, OSKAR-MORGENSTERN-PLATZ 1, 1090 WIEN (AUSTRIA)
Email address: `giacomo.sodini@univie.ac.at`