

The quadratic Wasserstein metric for inverse data matching

Björn Engquist* Kui Ren† Yunan Yang‡

Abstract

This work characterizes, analytically and numerically, two major effects of the quadratic Wasserstein (W_2) distance as the measure of data discrepancy in computational solutions of inverse problems. First, we show, in the infinite-dimensional setup, that the W_2 distance has a smoothing effect on the inversion process, making it robust against high-frequency noise in the data but leading to a reduced resolution for the reconstructed objects at a given noise level. Second, we demonstrate that, for some finite-dimensional problems, the W_2 distance leads to optimization problems that have better convexity than the classical L^2 and \mathcal{H}^{-1} distances, making it a more preferred distance to use when solving such inverse matching problems.

Key words. quadratic Wasserstein distance, inverse data matching, optimal transport, computational inverse problems, numerical optimization, preconditioning

AMS subject classifications 2010. 49N45, 49M99, 65K10, 65N21, 65M32, 74J25, 86A22

1 Introduction

This paper is concerned with inverse problems where we intend to match certain given data to data predicted by a (discrete or continuous) mathematical model, often called the forward model. To set up the problem, we denote by a function $m(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ ($d \geq 1$) the input of the mathematical model that we are interested in reconstructing from a given datum g . We denote by f the forward operator that maps the unknown quantity m to the datum g , that is

$$f(m) = g, \tag{1}$$

*Department of Mathematics and the Oden Institute, The University of Texas, Austin, TX 78712; engquist@ices.utexas.edu

†Department of Applied Physics and Applied Mathematics, Columbia University, New York, NY 10027; kr2002@columbia.edu

‡Courant Institute of Mathematical Sciences, New York University, New York, NY 10012; yunan.yang@nyu.edu

where the operator f is assumed to be nonlinear in general. We denote by $A := f'(m_0)$ the linearization of the operator f at the background m_0 . With a bit of abuse of notation, we write $Am = g$ to denote a linear inverse problem of the form (1) where $f(m) := Am$. The space of functions where we take our unknown object m , denoted by \mathcal{M} , and datum g , denoted by \mathcal{G} , as well as the exact form of the forward operator $f : \mathcal{M} \mapsto \mathcal{G}$, will be given later when we study specific problems.

Inverse problems for (1) are mostly solved computationally due to the lack of analytic inversion formulas. Numerical methods often reformulate the problem as a data matching problem where one takes the solution as the function m^* that minimizes the data discrepancy, measured in a metric \mathfrak{D} , between the model prediction $f(m)$ and the measured datum g . That is,

$$m^* = \arg \min_{m \in \mathcal{M}} \Phi(m), \quad \text{with,} \quad \Phi(m) := \frac{1}{2} \mathfrak{D}^2(f(m), g). \quad (2)$$

The most popular metric used in the past to measure the prediction-data discrepancy is the L^2 metric $\mathfrak{D}(f(m), g) := \|f(m) - g\|_{L^2(\mathbb{R}^d)}$ due to its mathematical and computational simplicity. Moreover, it is often the case that a regularization term is added to the mismatch functional $\Phi(m)$ to impose extra prior knowledge on the unknown m (besides of the fact that it is an element of \mathcal{M}) to be reconstructed.

In recent years, the quadratic Wasserstein metric [1, 30, 33] is proposed as an alternative for the L^2 metric in solving such inverse data matching problems [6, 7, 13, 18, 20, 19, 25, 26, 31, 35]. Numerical experiments suggest that the quadratic Wasserstein metric has attractive properties for some inverse data matching problems that the classical L^2 metric does not have [14, 36]. The objective of this work is trying to understand mathematically these numerical observations reported in the literature. More precisely, we attempt to characterize the numerical inversion of (1) based on the quadratic Wasserstein metric and compare that with the inversion based on the classical L^2 metric.

In the rest of the paper, we first briefly review some background materials on the quadratic Wasserstein metric and its connection to inverse data match problems in Section 2. We then study in Section 3 the Fourier domain behavior of the solutions to (1) in the asymptotic regime of the Wasserstein metric: the regime where the model prediction f and the datum g are sufficiently close. We show that in the asymptotic regime, the Wasserstein inverse solution tends to be smoother than the L^2 based inverse solution. We then show in Section 4 that this smoothing effect of the Wasserstein metric also exists in the non-asymptotic regime, but in a less explicit way. In Section 5, we demonstrate, in perhaps overly simplified settings, some advantages of the Wasserstein metric over the L^2 metric in solving some important inverse matching problems: inverse transportation, back-projection from (possibly partial) data, and deconvolution of highly concentrated sources. Numerical simulations are shown in Section 6 to demonstrate the main findings of our study. Concluding remarks are offered in Section 7.

2 Background and problem setup

Let f and g be two probability densities on \mathbb{R}^d that have the same total mass. The square of the quadratic Wasserstein distance between f and g denoted as $W_2^2(f, g)$, is defined as

$$W_2^2(f, g) := \inf_{\mathbf{T} \in \mathcal{T}} \int_{\mathbb{R}^d} |\mathbf{x} - \mathbf{T}(\mathbf{x})|^2 f(\mathbf{x}) d\mathbf{x}, \quad (3)$$

where \mathcal{T} is the set of measure-preserving maps from f to g . The map \mathbf{T} that achieves the infimum is called the optimal transport map between f and g . In the context of (1), the probability density f depends on the unknown function m . Therefore, $W_2^2(f, g)$ can be viewed as a mismatch functional of m for solving the inverse problem.

Since the quadratic Wasserstein distance is only defined between probability measures of the same total mass, one has to normalize f and g and turn them into probability densities when applying them to solve inverse matching problems where f and g cannot be interpreted as nonnegative probability density functions. This introduces new issues that need to be addressed [15]; see more discussions at the end of Section (7). The analysis in the rest of the paper assumes that f and g are both nonnegative and have the same total mass.

It is well-known by now that the quadratic Wasserstein distance between f and g is connected to a weighted $\dot{\mathcal{H}}^{-1}$ distance between them; see [33, Section 7.6] and [21, 28]. For any $s \in \mathbb{R}$, let $\mathcal{H}^s(\mathbb{R}^d)$ be the space of functions

$$\mathcal{H}^s(\mathbb{R}^d) := \{m(\mathbf{x}) : \|m\|_{\mathcal{H}^s(\mathbb{R}^d)}^2 := \int_{\mathbb{R}^d} \langle \boldsymbol{\xi} \rangle^{2s} |\widehat{m}(\boldsymbol{\xi})|^2 d\boldsymbol{\xi} < +\infty\}$$

where $\widehat{m}(\boldsymbol{\xi})$ denotes the Fourier transform of $m(\mathbf{x})$ and $\langle \boldsymbol{\xi} \rangle := \sqrt{1 + |\boldsymbol{\xi}|^2}$. When $s \geq 0$, $\mathcal{H}^s(\mathbb{R}^d)$ is the usual Hilbert space of functions with s square integrable derivatives, and $\mathcal{H}^0(\mathbb{R}^d) = L^2(\mathbb{R}^d)$. The space $\mathcal{H}^{-s}(\mathbb{R}^d)$ with $s > 0$ is understood as the dual of $\mathcal{H}^s(\mathbb{R}^d)$. We also introduce the space $\dot{\mathcal{H}}^s(\mathbb{R}^d)$, $s > 0$, with the (semi-) norm $\|\cdot\|_{\dot{\mathcal{H}}^s(\mathbb{R}^d)}$ defined through the relation

$$\|m\|_{\mathcal{H}^s(\mathbb{R}^d)}^2 = \|m\|_{L^2(\mathbb{R}^d)}^2 + \|m\|_{\dot{\mathcal{H}}^s(\mathbb{R}^d)}^2.$$

The space $\dot{\mathcal{H}}^{-s}(\mathbb{R}^d)$ is defined as the dual of $\dot{\mathcal{H}}^s(\mathbb{R}^d)$ via the norm

$$\|m\|_{\dot{\mathcal{H}}^{-s}} := \sup\{|\langle w, m \rangle| : \|w\|_{\dot{\mathcal{H}}^s} \leq 1\}. \quad (4)$$

It was shown [33, Section 7.6] that *asymptotically* W_2 is equivalent to $\dot{\mathcal{H}}_{(d\mu)}^{-1}$, where the subscript $(d\mu)$ indicates that the space is defined with respect to the reference probability measure $d\mu = f(\mathbf{x})d\mathbf{x}$. To be precise, if μ is the probability measure and $d\pi$ is an infinitesimal perturbation that has zero total mass, then

$$W_2(\mu, \mu + d\pi) = \|d\pi\|_{\dot{\mathcal{H}}_{(d\mu)}^{-1}} + o(d\pi). \quad (5)$$

This fact allows one to show that, for two positive probability measures μ and ν with densities f and g that are sufficiently regular, we have the following *non-asymptotic* equivalence between W_2 and $\dot{\mathcal{H}}_{(d\mu)}^{-1}$:

$$\mathfrak{c}_1 \|\mu - \nu\|_{\dot{\mathcal{H}}_{(d\mu)}^{-1}} \leq W_2(\mu, \nu) \leq \mathfrak{c}_2 \|\mu - \nu\|_{\dot{\mathcal{H}}_{(d\mu)}^{-1}}, \quad (6)$$

for some constants $\mathbf{c}_1 > 0$ and $\mathbf{c}_2 > 0$. The second inequality is generic with $\mathbf{c}_2 = 2$ [28, Theorem 1] while the first inequality, proved in [21, Proposition 2.8] and [28, Theorem 5] independently, requires further that f and g be bounded from above.

In the rest of this paper, we study numerical solutions to the inverse data matching problem for (1) under three different mismatch functionals:

$$\Phi_{\mathcal{H}^s}(m) \equiv \frac{1}{2} \|f(m) - g\|_{\mathcal{H}^s}^2 := \frac{1}{2} \int_{\mathbb{R}^d} \langle \boldsymbol{\xi} \rangle^{2s} |\widehat{f}(m)(\boldsymbol{\xi}) - \widehat{g}(\boldsymbol{\xi})|^2 d\boldsymbol{\xi}, \quad (7)$$

$$\Phi_{\mathcal{H}_{(d\mu)}^s}(m) \equiv \frac{1}{2} \|f(m) - g\|_{\mathcal{H}_{(d\mu)}^s}^2 := \frac{1}{2} \int_{\mathbb{R}^d} \left| \widehat{\omega} \otimes [\langle \boldsymbol{\xi} \rangle^s (\widehat{f}(m)(\boldsymbol{\xi}) - \widehat{g}(\boldsymbol{\xi}))] \right|^2 d\boldsymbol{\xi}, \quad (8)$$

where $\omega(\mathbf{x}) = 1/g(\mathbf{x})$, \otimes denotes the convolution operation, and

$$\Phi_{W_2}(m) \equiv \frac{1}{2} W_2^2(f(m), g) := \frac{1}{2} \inf_{\mathbf{T} \in \mathcal{T}} \int_{\mathbb{R}^d} |\mathbf{x} - \mathbf{T}(\mathbf{x})|^2 f(m(\mathbf{x})) d\mathbf{x}. \quad (9)$$

Our main goal is to analyze the differences between the Fourier contents of the inverse matching results, a main motivation for us to choose the Fourier domain definition of the \mathcal{H} norms. These norms allow us to systematically study: (i) the differences between the L^2 (i.e. the special case of $s = 0$ of $\Phi_{\mathcal{H}^s}(m)$) and the \mathcal{H}^s , with a positive or negative s , inversion results; (ii) the differences between \mathcal{H}^s and $\mathcal{H}_{(d\mu)}^s$ inversion results caused by the weight $d\mu$; and (iii) the similarities and differences between $\mathcal{H}_{(d\mu)}^s$ and W_2 inversion results. This is our roadmap toward better understandings of the differences between L^2 -based and W_2 -based inverse data matching.

Remark 2.1. *Note that since the $\dot{\mathcal{H}}^s$ norm is only a shift away from the corresponding \mathcal{H}^s norm in the Fourier representation, by replacing $\langle \boldsymbol{\xi} \rangle$ with $|\boldsymbol{\xi}|$, we do not introduce extra mismatch functionals for those (semi-) norms. We will, however, discuss $\dot{\mathcal{H}}^s$ inversions when we study the corresponding \mathcal{H}^s inversions.*

Remark 2.2. *In the definition of the $\mathcal{H}_{(d\mu)}^s$ objective function, we take the weight function $\omega(\mathbf{x}) = 1/g(\mathbf{x})$ such that $\|f - g\|_{\dot{\mathcal{H}}_{(d\mu)}^{-1}}^2$ is consistent with the linearization of $W_2^2(f, g)$ [33].*

We refer interested readers to [33, 21, 28] for technical discussions on the results in (5) and (6) (under more general settings than what we present here) that connect W_2 with $\dot{\mathcal{H}}_{(d\mu)}^{-1}$. For our purpose, these results say that: (i) in the asymptotic regime when two signals f and g , both being probability density functions, are sufficiently close to each other, their W_2 distance can be well approximated by their $\dot{\mathcal{H}}_{(d\mu)}^{-1}$ distance; and (ii) if $W_2(f, g) = 0$, then $\|f - g\|_{\dot{\mathcal{H}}_{(d\mu)}^{-1}} = 0$ and vice versa, that is, the exact matching solutions to the model (1), if exists, are global minimizers to both $\Phi_{W_2}(m)$ and $\Phi_{\dot{\mathcal{H}}_{(d\mu)}^{-1}}(m)$. However, let us emphasize that the non-asymptotic equivalence in (6) does NOT imply that the functional $\Phi_{W_2}(m)$ and $\Phi_{\dot{\mathcal{H}}_{(d\mu)}^{-1}}(m)$ (if we define one) have exactly the same optimization landscape. In fact, numerical evidences show that the two functionals have different optimization landscapes that are both quite different from that of the L^2 mismatch functional $\Phi_{L^2}(m) := \Phi_{\mathcal{H}^0}(m)$; see for instance Section (5) for analytical and numerical evidences.

3 Frequency responses in asymptotic regime

We first study the Fourier-domain behavior of the solutions to (1) obtained by minimizing the functionals we introduced in the previous section. At the solution, $f(m)$ and g are sufficiently close to each other. Therefore their W_2 distance can be replaced with their $\mathcal{H}_{(d\mu)}^{-1}$ distance according to (5). In the leading order, the W_2 solution is simply the $\mathcal{H}_{(d\mu)}^{-1}$ solution in this regime.

3.1 \mathcal{H}^s -based inverse matching for linear problems

Let us start with a linear inverse matching problem given by the model:

$$Am = g_\delta, \quad (10)$$

where g_δ denotes the datum g in (1) polluted by an additive noise introduced in the measuring process. The subscript δ is used to denote the size (in appropriate norms to be specified soon) of the noise, that is, the size of $g_\delta - g$. Besides, we assume that g_δ is still in the range of the operator A . When the model is viewed as the linearization of the nonlinear model (1), m should be regarded as the perturbation of the background m_0 . The model perturbation is also often denoted as δm . We assume that the linear operator A is diagonal in the Fourier domain, that is, it has the symbol,

$$\widehat{A}(\boldsymbol{\xi}) \sim \langle \boldsymbol{\xi} \rangle^{-\alpha}, \quad (11)$$

for some $\alpha \in \mathbb{R}$. *This assumption is to make some of the calculations more concise but is not essential as we will comment on later; see Remark 3.3.* When the exponent $\alpha > 0$, the operator A is “smoothing”, in the sense that it maps a given m to an output with better regularity than m itself. The inverse matching problem of solving for m in (10), on the other hand, is ill-conditioned (so would be the corresponding nonlinear inverse problem $f(m) = g$ if A is regarded as the linearization of f). The size of α , to some extent, can describe the degree of ill-conditionedness of the inverse matching problem.

We assume *a priori* that $m \in \mathcal{H}^\beta(\mathbb{R}^d)$ for some $\beta > 0$. Therefore, A could be viewed as an operator $A : \mathcal{H}^\beta \mapsto \mathcal{H}^{\beta+\alpha}$. We now look at the inversion of the problem under the \mathcal{H}^s ($s \leq \alpha + \beta$) framework.

We seek the solution of the inverse problem as the minimizer of the \mathcal{H}^s functional $\Phi_{\mathcal{H}^s}(m)$, defined as in (7) with $f(m) = Am$ and g replaced with g_δ . We verify that the Fréchet derivative of $\Phi_{\mathcal{H}^s} : \mathcal{H}^\beta \mapsto \mathbb{R}_{\geq 0}$ at m in the direction δm is given by

$$\Phi'_{\mathcal{H}^s}(m)[\delta m] = \int_{\mathbb{R}^d} \widehat{A}^*(\boldsymbol{\xi}) \left\{ \langle \boldsymbol{\xi} \rangle^{2s} \left[\widehat{A}(\boldsymbol{\xi}) \widehat{m}(\boldsymbol{\xi}) - \widehat{g}_\delta(\boldsymbol{\xi}) \right] \right\} \widehat{\delta m}(\boldsymbol{\xi}) d\boldsymbol{\xi},$$

where we used A^* to denote the L^2 adjoint of the operator A . The minimizer of $\Phi_{\mathcal{H}^s}$ is located at the place where its Fréchet derivative vanishes. Therefore the minimizer solves the following (modified) normal equation at frequency $\boldsymbol{\xi}$:

$$\widehat{A}^*(\boldsymbol{\xi}) \left\{ \langle \boldsymbol{\xi} \rangle^{2s} \widehat{A} \right\} \widehat{m} = \widehat{A}^*(\boldsymbol{\xi}) \left\{ \langle \boldsymbol{\xi} \rangle^{2s} \widehat{g}_\delta(\boldsymbol{\xi}) \right\}. \quad (12)$$

The solution at frequency $\boldsymbol{\xi}$ is therefore

$$\widehat{m}(\boldsymbol{\xi}) = \left(\widehat{A}^*(\boldsymbol{\xi}) (\langle \boldsymbol{\xi} \rangle^{2s} \widehat{A}) \right)^{-1} \widehat{A}^*(\boldsymbol{\xi}) \left(\langle \boldsymbol{\xi} \rangle^{2s} \widehat{g}_\delta(\boldsymbol{\xi}) \right).$$

We can then perform an inverse Fourier transform to find the solution in the physical space. The result is

$$m = \left(A^* P A \right)^{-1} A^* P g_\delta, \quad P := (\mathcal{I} - \Delta)^{s/2}, \quad (13)$$

where the operator $(\mathcal{I} - \Delta)^{s/2}$ is defined through the relation

$$(\mathcal{I} - \Delta)^{s/2} m := \mathcal{F}^{-1} \left(\langle \boldsymbol{\xi} \rangle^s \widehat{m} \right),$$

\mathcal{F}^{-1} being the inverse Fourier transform, Δ being the usual Laplacian operator, and \mathcal{I} being the identity operator.

Key observations. Let us first remark that the calculations above can be carried out in the same way if the \mathcal{H}^s norm is replaced with the $\dot{\mathcal{H}}^s$ norm. The only changes are that $\langle \boldsymbol{\xi} \rangle$ should be replaced with $|\boldsymbol{\xi}|$ and the operator $\mathcal{I} - \Delta$ in P has to be replaced with $-\Delta$.

When $s = 0$, assuming that $\alpha + \beta \geq 0$, the above solution reduces to the classical L^2 least-squares solution $m = (A^* A)^{-1} A^* g_\delta$. Moreover, when A is invertible (so will be A^*), the solution can be simplified to $m = A^{-1} g_\delta$, using the fact that $(A^* P A)^{-1} = A^{-1} P^{-1} A^{-*}$, which is simply the true solution to the original problem (10). Therefore, in the same manner, as the classical L^2 least-squares method, the least-squares method based on the \mathcal{H}^s norm does not change the solution to the original inverse problem when it is uniquely solvable. This is, however, not the case for the $\dot{\mathcal{H}}^s$ inversion in general. For instance, $\dot{\mathcal{H}}^1$ inversion only matches the derivative of the predicted data to the measured data.

When $s > 0$, P is a differential operator. Applying P to the datum g_δ amplifies high-frequency components of the datum. When $s < 0$, P is a (smoothing) integral operator. Applying P to the datum g_δ suppresses high-frequency components of the datum. Even though the presence of the operator P in $(A^* P A)^{-1}$ will un-do the effect of P on the datum in a perfect world (when A is invertible, and all calculations are done with arbitrary precision), when operated under a given accuracy, inversion with $s < 0$ is less sensitive to high-frequency noise in the data while inversion with $s > 0$ is more sensitive to high-frequency noise in the data, compared to the case of $s = 0$ (that is, the classical L^2 least-squares inversion). Therefore, inversion with $s \neq 0$ can be seen as a “preconditioned” (by the operator P) L^2 least-squares inversion.

3.2 Resolution analysis of linear inverse matching

We now perform a simple resolution analysis, following the calculations in [3], on the \mathcal{H}^s inverse matching result for the linear model (10).

Theorem 3.1. Let A be given as in (11) and R_c an approximation to A^{-1} defined through its symbol:

$$\widehat{R}_c(\boldsymbol{\xi}) \sim \begin{cases} \langle \boldsymbol{\xi} \rangle^\alpha, & |\boldsymbol{\xi}| < \xi_c \\ 0, & |\boldsymbol{\xi}| > \xi_c \end{cases}.$$

Let $\delta = \|g_\delta - g\|_{\mathcal{H}^s}$ be the \mathcal{H}^s norm of the additive noise in g_δ . Then the reconstruction error $\|m - m_\delta^c\|_{L^2}$, with $m_\delta^c := R_c g_\delta$ obtained as the minimizer of $\Phi(m)_{\mathcal{H}^s}$, is bounded by

$$\|m - m_\delta^c\|_{L^2} \lesssim \|m\|_{\mathcal{H}^\beta}^{\frac{\alpha-s}{\alpha+\beta-s}} \delta^{\frac{\beta}{\alpha+\beta-s}}. \quad (14)$$

This optimal bound is achieved when we select

$$\langle \xi_c \rangle^{-1} \sim (\delta \|m\|_{\mathcal{H}^\beta}^{-1})^{\frac{1}{\alpha+\beta-s}}. \quad (15)$$

Proof. Following classical results in [12], it is straightforward to verify that the L^2 difference between the true solution m and the approximated noisy solution m_δ^c is

$$\begin{aligned} \|m - m_\delta^c\|_{L^2} &= \|m - R_c g_\delta\|_{L^2} = \|m - R_c g + R_c g - R_c g_\delta\|_{L^2} \\ &= \|(\mathcal{I} - R_c A)m + R_c(g - g_\delta)\|_{L^2} \leq \|(\mathcal{I} - R_c A)m\|_{L^2} + \|R_c(g - g_\delta)\|_{L^2}. \end{aligned} \quad (16)$$

We then verify that operators $\mathcal{I} - R_c A : \mathcal{H}^\beta(\mathbb{R}^d) \mapsto L^2(\mathbb{R}^d)$ and $R_c : \mathcal{H}^s(\mathbb{R}^d) \mapsto L^2(\mathbb{R}^d)$ have the following norms respectively

$$\|R_c\|_{\mathcal{H}^s \mapsto L^2} \sim \langle \xi_c \rangle^{\alpha-s} \quad \text{and} \quad \|(\mathcal{I} - R_c A)\|_{\mathcal{H}^\beta \mapsto L^2} \sim \langle \xi_c \rangle^{-\beta}.$$

This allows us to conclude that

$$\|m - m_\delta^c\|_{L^2} \leq \|R_c\|_{\mathcal{H}^s \mapsto L^2} \delta + \|(\mathcal{I} - R_c A)\|_{\mathcal{H}^\beta \mapsto L^2} \|m\|_{\mathcal{H}^\beta} \lesssim \langle \xi_c \rangle^{\alpha-s} \delta + \langle \xi_c \rangle^{-\beta} \|m\|_{\mathcal{H}^\beta}.$$

We can now select $\langle \xi_c \rangle \sim (\delta^{-1} \|m\|_{\mathcal{H}^\beta})^{\frac{1}{\alpha+\beta-s}}$, i.e. the relation given in (15), to minimize the error of the reconstruction, which gives the bound in (14). \square

Optimal resolution. One main message carried by the theorem above is that reconstruction based on the \mathcal{H}^s mismatch has a spatial resolution

$$\varepsilon := \langle \xi_c \rangle^{-1} \sim \delta^{\frac{1}{\alpha+\beta-s}},$$

under the conditions in the theorem. At a fixed noise level δ , for fixed α and β , the optimal resolution of the inverse matching result degenerates when s gets smaller. The case of $s = 0$ corresponds to the usual reconstruction in the L^2 framework. The optimal resolution one could get in this case is decided by $\delta^{\frac{1}{\alpha+\beta}}$. When $0 < s (\leq \alpha + \beta)$, the best resolution one could get is better than the L^2 case in a perfect world. When $s < 0$, the reconstructions in the \mathcal{H}^s framework provides an optimal resolution that is worse than the L^2 case. In other words, the reconstructions based on the negative norms appear to be smoother than optimal L^2 reconstructions in this case. See Section 6 for numerical examples that illustrate this resolution analysis.

However, we should emphasize that the above simple calculation only provides the best-case scenarios. It does not tell us exactly how to achieve the best results in a general setup (when the symbol of A , i.e., the singular value decomposition of A in the discrete case, is not precisely known). Nevertheless, the guiding principle of the analysis is well demonstrated: least-squares with a stronger (than the L^2) norm yield higher resolution reconstructions while least-squares with a weaker norm yield lower (again compared to the L^2 case) resolution reconstructions in the best case.

3.3 $\mathcal{H}_{(d\mu)}^s$ -based inverse matching for linear problems

Inverse matching with the weighted \mathcal{H}^s norm can be analyzed in the same manner to study the impact of the weight on the inverse matching result. The solution m to (10) in this case is sought as the minimizer of the functional $\Phi_{\mathcal{H}_{(d\mu)}^s}(m)$ defined in (8) with $f(m) = Am$ and $g = g_\delta$. This means that the weight $\omega = 1/g_\delta$ in our definition of the objective function.

Following the same calculation as in the previous subsection, we find that the minimizer of the functional $\Phi_{\mathcal{H}_{(d\mu)}^s}(m)$ solves the following normal equation at frequency ξ :

$$\widehat{B}^* \widehat{\omega} \circledast (\langle \xi \rangle^s \widehat{A} \widehat{m}) = \widehat{B}^* \widehat{\omega} \circledast (\langle \xi \rangle^s \widehat{g}_\delta) \quad (17)$$

where \widehat{B}^* is the L^2 adjoint of the operator \widehat{B} defined through the relation $\widehat{B} \widehat{m} := \widehat{\omega} \circledast (\langle \xi \rangle^s \widehat{A} \widehat{m})$.

We first observe that the right-hand side of (17) and that of (12) are different. In (12), the ξ -th Fourier mode of the datum g_δ is amplified or attenuated, depending on the sign of s , by a factor of $\langle \xi \rangle^{2s}$. While in (17), this mode is further convoluted with other modes of the datum after the amplification/attenuation. The convolution induces mixing between different modes of the datum. Therefore, inverse matching with the weighted norm cannot be done mode by mode as what we did for the unweighted norm, even when we assume that the forward operator A is diagonal. However, main effect of the norm on the inversion, the smoothing/sharpening effect introduced by the $\langle \xi \rangle^{2s}$ factor (half of which come from the factor $\langle \xi \rangle^s$ in front of \widehat{g}_δ while the other half come from the factor $\langle \xi \rangle^s$ in \widehat{B}), are the same in both the unweighted \mathcal{H}^s and the weighted $\mathcal{H}_{(d\mu)}^s$ norms.

The inverse matching solution, in this case, written in physical space, is:

$$m = \left(A^* P_g A \right)^{-1} A^* P_g g_\delta, \quad P_g := (\mathcal{I} - \Delta)^{s/2} \omega (\mathcal{I} - \Delta)^{s/2}. \quad (18)$$

We can again compare this with the unweighted solution in (13). The only difference is the introduction of the inhomogeneity, which depends on the datum g_δ , in the preconditioning operator P by replacing it with P_g . When $0 < s (\leq \alpha + \beta)$, P and P_g are (local) differential operators. Roughly speaking, compared to P , P_g emphasizes places where g_δ is small, be reminded that $\omega = 1/g_\delta$, or the s -th order derivative of g_δ is large. At those locations, P_g amplifies the same modes of the datum g_δ more than P does. When $s < 0$, P and P_g are non-local operators. The impact of g_δ is more global (as we have seen in the previous paragraph

in the Fourier domain). It is hard to precisely characterize the impact of g_δ without knowing its form explicitly. However, we can still see, for instance, from (17), that the smoother g_δ is, the smoother the inverse matching result will be (since \widehat{g}_δ has fast decay and the convolution will further smooth out $\langle \xi \rangle^s g_\delta$). If g_δ is very rough, say that it behaves like Gaussian noise, then \widehat{g}_δ decays very slowly. The convolution, in this case, will not smooth out $\langle \xi \rangle^s g_\delta$ as much as in the previous case. The main effect of $\mathcal{H}_{(d\mu)}^s$ on the inverse matching result in this case mainly comes from the norm, not the weight.

Remark 3.2. *Thanks to the asymptotic equivalence between $\dot{\mathcal{H}}_{(d\mu)}^{-1}$ and W_2 in (5), the smoothing effect we observe in this section for the $\mathcal{H}_{(d\mu)}^{-1}$ inverse matching (and therefore $\dot{\mathcal{H}}_{(d\mu)}^{-1}$ inverse matching since $\dot{\mathcal{H}}_{(d\mu)}^{-1}$ is only different from $\mathcal{H}_{(d\mu)}^{-1}$ on the zeroth moment in the Fourier domain) is also what we observe in the W_2 inverse matching. This observation will be demonstrated in more detail in our numerical simulations in Section 6.*

3.4 Iterative solution of nonlinear inverse matching

The simple analysis in the previous sections based on the linearized quadratic Wasserstein metric, i.e., a weighted $\dot{\mathcal{H}}^{-1}$ norm, on the inverse matching of linear model (10) does not translate directly to the case of inverse matching with the fully nonlinear model (1). Nevertheless, the analysis does provide us some insights.

Let us consider an iterative matching algorithm for the nonlinear problem, starting with a given initial guess m_0 , characterized by the following iteration:

$$m_{k+1} = m_k + \ell_k \zeta_k, \quad k \geq 0, \quad (19)$$

where ζ_k is a chosen descent direction of the objective functional at iteration k , and ℓ_k is the step length at this iteration. For simplicity, let us take the steepest descent method where the descent direction is taken as the negative gradient direction. Following the calculations in the previous section, we verify that the Fréchet derivative of $\Phi_{\mathcal{H}_{(d\mu)}^s}(m) : \mathcal{H}^\beta \mapsto \mathbb{R}_{\geq 0}$ at the current iteration m_k in the direction δm is given by

$$\Phi'_{\mathcal{H}_{(d\mu)}^s}(m_k)[\delta m] = \int_{\mathbb{R}^d} \widehat{\omega} \otimes [\langle \xi \rangle^s (\widehat{f(m_k)} - \widehat{g_\delta})] \widehat{\omega} \otimes \langle \xi \rangle^s f'(\widehat{m_k})[\delta m] d\xi, \quad (20)$$

assuming that the forward model $f : \mathcal{H}^\beta \mapsto \mathcal{H}_{(d\mu)}^s$ is Fréchet differentiable at m_k with derivative $f'(m_k)[\delta m]$. This leads to the following descent direction ζ_k chosen by a gradient descent method:

$$\zeta_k = - \left(f'(m_k)^* P_g f'(m_k) \right)^{-1} f'(m_k)^* P_g (f(m_k) - g_\delta), \quad (21)$$

Let us compare this with the descent direction resulted from the L^2 least-squares functional:

$$\zeta_k = \left(f'(m_k)^* f'(m_k) \right)^{-1} f'(m_k)^* (f(m_k) - g). \quad (22)$$

We see that the iterative process of the $\mathcal{H}_{(d\mu)}^s$ inverse matching can be viewed as a preconditioned version of the corresponding L^2 iteration. The preconditioning operator, P_g , depends on the datum g_δ but is independent of the iteration. When the iteration is stopped after a finite step, the effect we observed for linear problems, that is, the smoothing effect of P_g in the case of $s < 0$ or its de-smoothing effect in the case of $s > 0$, is carried to the solution of nonlinear problems.

Wasserstein smoothing in the asymptotic regime. To summarize, when the model predictions and the measured data are sufficiently close to each other, inverse matching with the quadratic Wasserstein metric, or equivalently the $\dot{\mathcal{H}}_{(d\mu)}^{-1}$ metric, can be viewed as a preconditioned L^2 -based inverse matching. The preconditioning operator is roughly the inverse Laplacian operator with a coefficient given by the datum. The optimal resolution of the inversion result from the Wasserstein metric, with data at a given noise level δ is roughly of the order $\delta^{\frac{1}{\alpha+\beta+1}}$ (α being the order of the operator $f'(m)$ at the optimal solution and $m \in \mathcal{H}^\beta$) instead of $\delta^{\frac{1}{\alpha+\beta}}$ as given in the L^2 case. The shape of the datum g_δ distorts the Wasserstein matching result slightly from the inverse matching result with a regular $\dot{\mathcal{H}}^{-1}$ (semi-) norm.

Remark 3.3. *The assumption that the linear operator A is diagonal in the Fourier domain, given in (11), simplifies the calculations in this section. The assumption is not necessary at all to show the preconditioning effect of the W_2 metric. Without this assumption, we need to replace all the multiplication of \hat{A} in the Fourier domain with convolutions. The final results remain the same. The assumption is indeed necessary in order to write down the approximate inverse operator R_c in the Fourier domain explicitly. This leads to a precise resolution characterization in Theorem 3.1 for this reconstruction operator.*

4 Wasserstein iterations in non-asymptotic regime

As we have seen from the previous sections, in the asymptotic regime, the smoothing effect of the quadratic Wasserstein metric in solving inverse matching problems can be characterized relatively precise thanks to the equivalence between W_2 and $\dot{\mathcal{H}}_{(d\mu)}^{-1}$ given in (5). The demonstrated smoothing effect makes W_2 -based inverse matching very robust against high-frequency noise in the measured data. This phenomenon has been reported in the numerical results published in recent years [6, 13, 14, 35, 36] and is one of the main reasons that W_2 is considered as a good alternative for L^2 -based matching methods. In this section, we argue that the smoothing effect of W_2 can also be observed in the non-asymptotic regime, that is, a regime where signals f and g are sufficiently far away from each other. *The smoothing effect in the non-asymptotic regime implies that the landscape of the W_2 objective functional is smoother than that of the classical L^2 objective functional.*

To see the smoothing effect of W_2 in non-asymptotic regime, we analyze the inverse matching procedure described by the iterative scheme (19) for the objective functional $\Phi_{W_2}(m)$, defined in (9). For the sake of being technically correct, we assume that the

data we are comparing in this section are sufficiently regular. More precisely, we assume that $f \in \mathcal{C}^{0,\alpha}(\mathbb{R}^d)$ and $g \in \mathcal{C}^{0,\alpha}(\mathbb{R}^d)$ for some $\alpha > 0$. We also assume that the map $m \mapsto f(\mathcal{H}^\beta \rightarrow \mathcal{C}^{0,\alpha})$ is Fréchet differentiable at any admissible m and denote by $f'(m)[\delta m]$ the derivative in direction δm . We can then write down the variation of $\Phi_{W_2}(m) : \mathcal{H}^\beta \mapsto \mathbb{R}_{\geq 0}$ at the current iteration m_k in the direction δm , following the differentiability result of $W_2^2(f, g)$ with respect to f along mass preserving perturbations [33, Theorem 8.13]. It is,

$$\Phi'_{W_2}(m_k)[\delta m] = \int_{\mathbb{R}^d} \left(\frac{|\mathbf{x} - \mathbf{T}_k(\mathbf{x})|^2}{2} f'(m_k)[\delta m] - (\mathbf{x} - \mathbf{T}_k(\mathbf{x}))f(\mathbf{x}) \cdot \mathbf{T}'_k[f'(m_k)[\delta m]] \right) d\mathbf{x}, \quad (23)$$

where \mathbf{T}_k denotes the optimal transport map at iteration k (that is, for m_k), and $\mathbf{T}'_k[\delta f]$ denotes the variation of \mathbf{T}_k with respect to f (not m) in the direction δf . We emphasize again that δm is selected such that $\int_{\mathbb{R}^d} f(m_k) d\mathbf{x} = \int_{\mathbb{R}^d} f(m_k + \delta m) d\mathbf{x}$ which is necessary since the space of probability densities with the W_2 metric is not a linear space.

Following the optimal transport theorem of Brenier [33], the optimal transport map at the current iteration k , \mathbf{T}_k , is given as $\mathbf{T}_k(\mathbf{x}) := \nabla u(\mathbf{x})$ where u is the unique (up to a constant) convex solution to the Monge-Ampère equation:

$$\det(D^2 u(\mathbf{x})) = f(m_k(\mathbf{x}))/g(\nabla u(\mathbf{x})), \quad u \text{ being convex.} \quad (24)$$

Here D^2 is the Hessian operator defined through the Hessian matrix $D^2 u := (\frac{\partial^2 u}{\partial x_i \partial x_j})$ (with the notation $\mathbf{x} = (x_1, \dots, x_d)$). Interested readers are referred to [33] and the references therein for more detailed mathematical study on the existence and uniqueness of solutions to the Monge-Ampère equation (24).

Let $\varphi := u'(f(m_k))[\delta f]$ be the Fréchet derivative of u at $f(m_k)$ in the direction δf , we then verify that φ solves the following second-order elliptic equation to the leading order:

$$\sum_{ij} a_{ij} \frac{\partial^2 \varphi}{\partial x_i \partial x_j} + \sum_j b_j \frac{\partial \varphi}{\partial x_j} = \delta f, \quad (25)$$

where $b_j = \det(D^2 u) \partial_{x_j} g(\mathbf{T}_k(\mathbf{x}))$ while a_{ij} depend on the dimension. When $d = 2$, $a_{ij} = -g(\mathbf{T}_k(\mathbf{x})) \frac{\partial^2 u}{\partial x_j \partial x_i}$ ($i \neq j$) and $a_{ii} = g(\mathbf{T}_k(\mathbf{x})) \frac{\partial^2 u}{\partial x_j \partial x_j}$ ($i \neq j$). When $d = 3$, we have

$$a_{ij} = g(\mathbf{T}_k(\mathbf{x})) \begin{cases} \frac{\partial^2 u}{\partial x_k \partial x_i} \frac{\partial^2 u}{\partial x_j \partial x_k} - \frac{\partial^2 u}{\partial x_j \partial x_i} \frac{\partial^2 u}{\partial x_k \partial x_k}, & i \neq j \neq k, \\ \frac{\partial^2 u}{\partial x_{k'} \partial x_{k'}} \frac{\partial^2 u}{\partial x_k \partial x_k} - \frac{\partial^2 u}{\partial x_{k'} \partial x_k} \frac{\partial^2 u}{\partial x_k \partial x_{k'}}, & i = j \neq k \neq k'. \end{cases}$$

Let ψ be the solution to the (adjoint) equation:

$$\sum_{ij} a_{ij} \frac{\partial^2 \psi}{\partial x_i \partial x_j} - \sum_j b_j \frac{\partial \psi}{\partial x_j} = -\nabla \cdot \left((\mathbf{x} - \mathbf{T}_k(\mathbf{x}))f(\mathbf{x}) \right). \quad (26)$$

It is then straightforward to verify, following standard adjoint state calculations [34], that update direction can be written as

$$\zeta_k(\mathbf{x}) = f'^*(m_k) \left[\frac{|\mathbf{x} - \mathbf{T}_k(\mathbf{x})|^2}{2} + \psi(\mathbf{x}) \right], \quad (27)$$

where $f'^*(m_k)$ denotes the L^2 adjoint of the operator $f'(m_k)$.

We first observe that unlike in the classical L^2 case where $f'^*(m_k)$ is applied directly to the residual $f(m_k) - g$, that is, $\zeta_k(\mathbf{x}) = f'^*(m_k) [f(m_k) - g]$, the descent direction here depends on the model prediction $f(m_k)$ and the datum g only implicitly through the transfer map $\mathbf{I} - \mathbf{T}_k$ and its variation with respect to m . Only in the asymptotic regime of g being very close to f can we make the connection between $\frac{|\mathbf{x} - \mathbf{T}_k(\mathbf{x})|^2}{2} + \psi(\mathbf{x})$ and the normalized residual. This is where the $\mathcal{H}_{(d\mu)}^{-1}$ approximation to W_2 comes from.

From Caffarelli's regularity theory (c.g. [33, Theorem 4.14]), which states that when $f \in \mathcal{C}^{0,\alpha}(\mathbb{R}^d)$ and $g \in \mathcal{C}^{0,\alpha}(\mathbb{R}^d)$ we have that the Monge-Ampère solution $u \in \mathcal{C}^{2,\alpha}(\mathbb{R}^d)$, we see that $(\mathbf{x} - \mathbf{T}_k(\mathbf{x}))$ is at least $\mathcal{C}^{1,\alpha}$. Therefore the solution to the adjoint problem, ψ , is in $\mathcal{C}^{2,\alpha}$ by standard theory for elliptic partial differential equations when the problem is not degenerate and in $\mathcal{C}^{1,\alpha}$ if it is degenerate. Therefore, $\frac{|\mathbf{x} - \mathbf{T}_k(\mathbf{x})|^2}{2} + \psi(\mathbf{x}) \in \mathcal{C}^{1,\alpha}$ is one derivative smoother than f and g (and therefore the residual). This is exactly what the preconditioning operator P (with $s = -1$) did to the residual in the asymptotic regime, for instance, as shown in (13). This shows that W_2 inverse matching has smoothing effect even in the non-asymptotic regime.

In one-dimensional case, we can see the smoothing effect more explicitly since we are allowed to construct the optimal map explicitly in this case. Let F and G be the cumulative density functions for f and g respectively. The optimal transportation theorem in one-dimensional setting (c.g. [33, Theorem 2.18]) then says that the optimal transportation map from f to g is given by $T(x) = G^{-1} \circ F(x)$. This allows us to verify that, the gradient of $\Phi_{W_2}(m)$ at m_k in direction δm , given in (23), is simplified to:

$$\begin{aligned} \Phi'_{W_2}(m_k)[\delta m] &= \int_{\mathbb{R}} \left(\frac{(x - T_k(x))^2}{2} + p_k(+\infty) - p_k(x) \right) f'(m_k)[\delta m] dx \\ &= \int_{\mathbb{R}} f'^*(m_k) \left[\frac{(x - T_k(x))^2}{2} - p_k(+\infty) + p_k(x) \right] \delta m(x) dx \end{aligned} \quad (28)$$

where the function $p_k(x)$ is defined as $p_k(x) = \int_{-\infty}^x \frac{(y - T_k(y))f(m_k(y))}{g(T_k(y))} dy$. Therefore the descent direction (27) simplifies to

$$\zeta_k(\mathbf{x}) = f'^*(m_k) \left[\frac{(x - T_k(x))^2}{2} - p_k(+\infty) + p_k(x) \right]. \quad (29)$$

It is clear from (29) that the gradient of $\Phi_{W_2}(m)$ at iteration k depends only on the anti-derivatives of $f(m_k)$, g and $f(m_k)/g(T_k)$, through $T_k(x)$ and p_k . Therefore, it is smoother

than the Fréchet derivative of $\Phi_{L^2}(m)$ in general, whether or not the signals $f(m_k)$ and g are close to each other. This shows why the smoothing effect of W_2 exists also in non-asymptotic regime.




Figure 1: The gradients of the objective functions $\Phi_{L^2}(m)$ (left), $\Phi_{\mathcal{H}^{-1}}(m)$ (middle) and $\Phi_{W_2}(m)$ (right) at the initial guess for the inverse diffusion problem in Section 6.2 in the one-dimensional domain $\Omega = (0, 1)$.




Figure 2: The gradients of the objective functions $\Phi_{L^2}(m)$ (left), $\Phi_{\mathcal{H}^{-1}}(m)$ (middle) and $\Phi_{W_2}(m)$ (right) at the initial guess for the inverse diffusion problem in Section 6.2 in the two-dimensional domain $\Omega = (0, 1) \times (0, 1)$.

To provide some numerical evidences, we show in Figure 1 and Figure 2 some gradients of the L^2 and W_2 objective functions, with respect to the absorption coefficient σ (i.e. $m = \sigma$), for the inverse diffusion problem we study in Section 6.2, in one- and two-dimensional domains $\Omega = (0, 1)$ and $\Omega = (0, 1) \times (0, 1)$ respectively. The synthetic data, generated by applying the forward operator to the true absorption coefficient and then adding multiplicative random noise, contains roughly 5% of random noise. We intentionally choose initial guesses to be relatively far away from the true coefficient so that the model prediction $f(m)$ is far from the data g to be matched. We are not interested in a direct quantitative comparison between the gradient of the Wasserstein objective function and that of the L^2 objective function since we do not have a good basis for the comparison. However, it is evident

from these numerical results that the gradient of the Wasserstein functional is smoother, or contains fewer frequencies to be more precise, compared to the corresponding L^2 case.

5 Wasserstein inverse matching for transportation and convolution

Its robustness against high-frequency noise in data, resulted from its smoothing effect we demonstrated in the previous two sections, is not the only reason why W_2 is thought as better than L^2 for many inverse data matching problems. We show in this section another advantage of the W_2 distance in studying inverse matching problems: its convexity with respect to translation and dilation of signals.

5.1 W_2 convexity with respect to affine transformations

For a given probability density function ϕ with finite moments $\mathbf{M}_1 := \int_{\mathbb{R}^d} \mathbf{x}\phi(\mathbf{x})d\mathbf{x}$ and $\mathbf{M}_2 := \int_{\mathbb{R}^d} |\mathbf{x}|^2\phi(\mathbf{x})d\mathbf{x}$, we define:

$$f(m(\mathbf{x})) := \frac{1}{\sqrt{|\Sigma|}}\phi(\Sigma^{-\frac{1}{2}}(\mathbf{x} - \lambda\bar{\mathbf{x}})). \quad (30)$$

where $m := (\Sigma, \lambda, \bar{\mathbf{x}})$ with $\Sigma \in \mathbb{R}^{d \times d}$ symmetric and positive-definite, $\lambda \in \mathbb{R}$ and $\bar{\mathbf{x}} \in \mathbb{R}^d$. This f is simply a translation (by $\lambda\bar{\mathbf{x}}$) and dilation (by $\Sigma^{1/2}$) of the function ϕ . We verify that $\int_{\mathbb{R}^d} f(m(\mathbf{x}))d\mathbf{x} = \int_{\mathbb{R}^d} \phi(\mathbf{x})d\mathbf{x}$.

Let $g = f(m_g)$ be generated from $m_g := (\Sigma_g, \lambda_g, \bar{\mathbf{x}}_g)$ with $\Sigma_g \in \mathbb{R}^{d \times d}$ symmetric and positive-definite. Then we check that the optimal transport map from f to g is given by $\mathbf{T}(\mathbf{x}) = \Sigma_g^{1/2}\Sigma^{-1/2}(\mathbf{x} - \lambda\bar{\mathbf{x}}) + \lambda_g\bar{\mathbf{x}}_g$. In other words, the function $u(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \bar{\mathbf{x}})^t\Sigma_g^{-1/2}\Sigma^{1/2}(\mathbf{x} - \bar{\mathbf{x}}) + \bar{\mathbf{x}}_g \cdot \mathbf{x}$ is a convex solution to the Monge-Ampère equation (24) with this (f, g) pair. This observation allows us to find that,

$$W_2^2(f, g) = |\lambda\bar{\mathbf{x}} - \lambda_g\bar{\mathbf{x}}_g|^2 + 2(\lambda\bar{\mathbf{x}} - \lambda_g\bar{\mathbf{x}}_g)^t(\Sigma^{1/2} - \Sigma_g^{1/2})\mathbf{M}_1 + \int_{\mathbb{R}^d} |(\Sigma^{1/2} - \Sigma_g^{1/2})\mathbf{x}|^2\phi(\mathbf{x})d\mathbf{x}. \quad (31)$$

This calculation shows that $W_2^2(f, g)$ is convex with respect to $\lambda\bar{\mathbf{x}} - \lambda_g\bar{\mathbf{x}}_g$ and $\Sigma^{1/2} - \Sigma_g^{1/2}$ for rather general probability density function ϕ .

For the purpose of comparison, we recall that the L^2 distance between f and g in this case is given by

$$\|f - g\|_{L^2}^2 = \left(\frac{1}{|\Sigma^{1/2}|} + \frac{1}{|\Sigma_g^{1/2}|}\right) \int_{\mathbb{R}^d} \phi^2(\mathbf{y})d\mathbf{y} - \frac{2}{|\Sigma_g^{1/2}|} \int_{\mathbb{R}^d} \phi(\mathbf{y})\phi(\Sigma_g^{-1/2}\Sigma^{1/2}\mathbf{y} + \Sigma_g^{-1/2}(\lambda\bar{\mathbf{x}} - \lambda_g\bar{\mathbf{x}}_g))d\mathbf{y}, \quad (32)$$

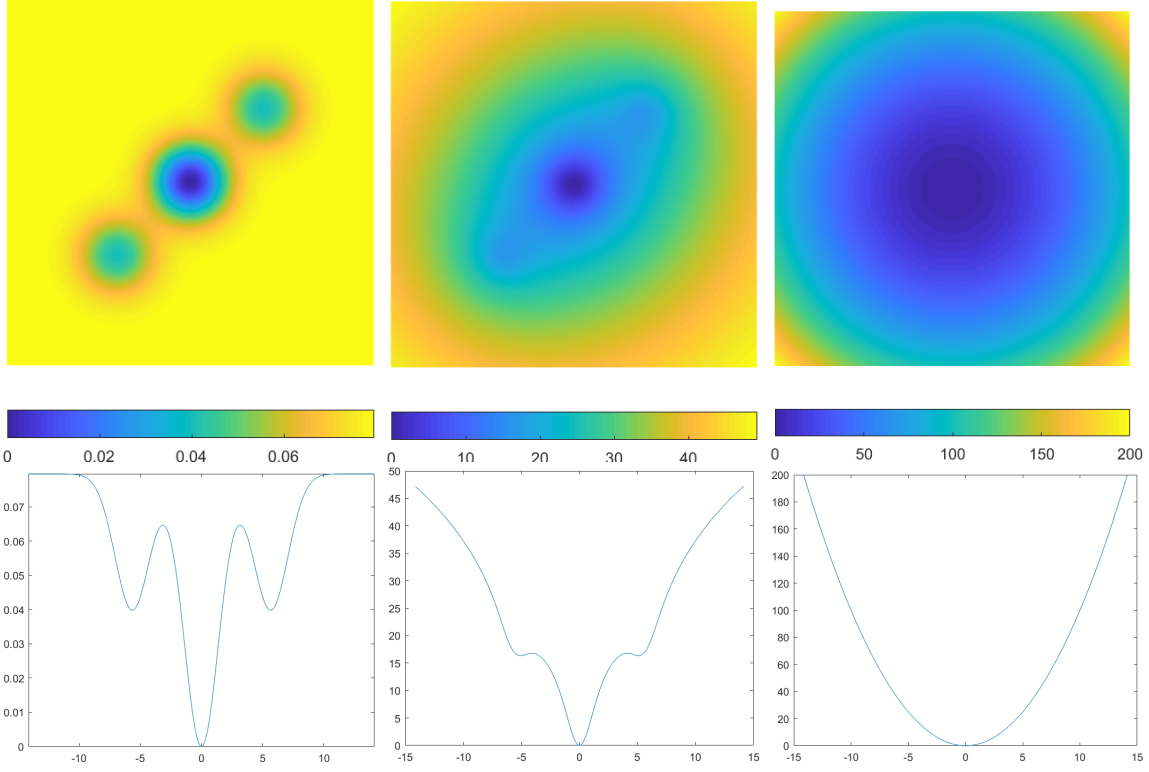


Figure 3: Top row: plots of the $\|f - g\|_{L^2}^2$ (left), $\|f - g\|_{\mathcal{H}^{-1}}^2$ (middle) and $W_2^2(f, g)$ (right) for the $\phi(\mathbf{x})$ given in (34) in the two-dimensional region $\lambda\bar{\mathbf{x}} - \lambda_g\bar{\mathbf{x}}_g \in [-10, 10]^2$ with $\Sigma = \Sigma_g = \mathbf{I}_2$. Bottom row: the corresponding cross-sections along the left bottom to top right diagonal.

and the \mathcal{H}^{-1} distance between f and g in this case is given by

$$\|f - g\|_{\mathcal{H}^{-1}}^2 = \int_{\mathbb{R}^d} |\boldsymbol{\xi}|^{-2} \left(|\widehat{\phi}(\Sigma^{1/2}\boldsymbol{\xi})|^2 + |\widehat{\phi}(\Sigma_g^{1/2}\boldsymbol{\xi})|^2 - 2\Re[\widehat{\phi}(\Sigma^{1/2}\boldsymbol{\xi})\overline{\widehat{\phi}(\Sigma_g^{1/2}\boldsymbol{\xi})}e^{i(\lambda\bar{\mathbf{x}} - \lambda_g\bar{\mathbf{x}}_g)\cdot\boldsymbol{\xi}}] \right) d\boldsymbol{\xi}. \quad (33)$$

By simple Taylor expansions, we see that $\|f - g\|_{L^2}^2$ and $\|f - g\|_{\mathcal{H}^{-1}}^2$ are indeed convex when $\Sigma^{1/2} - \Sigma_g^{1/2}$ and $\lambda\bar{\mathbf{x}} - \lambda_g\bar{\mathbf{x}}_g$ are sufficiently small. However, neither $\|f - g\|_{L^2}^2$ nor $\|f - g\|_{\mathcal{H}^{-1}}^2$ is globally convex with respect to $\Sigma^{1/2} - \Sigma_g^{1/2}$ or $\lambda\bar{\mathbf{x}} - \lambda_g\bar{\mathbf{x}}_g$ for a general probability density $\phi(\mathbf{x})$. This is one of the major differences between W_2 , L^2 and \mathcal{H}^{-1} . To visualize this advantage of W_2 , we plot in Figure 3 the functionals $W_2^2(f, g)$, $\|f - g\|_{L^2}^2$ and $\|f - g\|_{\mathcal{H}^{-1}}^2$ in the two-dimensional region $\Omega = [-10, 10]^2$ with fixed $\Sigma = \Sigma_g = \mathbf{I}_2$ for

$$\phi(\mathbf{x}) = \frac{1}{4\pi} \left[e^{-\frac{1}{2}|\mathbf{x}-\mathbf{y}_1|^2} + e^{-\frac{1}{2}|\mathbf{x}-\mathbf{y}_2|^2} \right], \quad \mathbf{y}_1 = (-2, -2), \quad \mathbf{y}_2 = (2, 2). \quad (34)$$

The plots clearly demonstrate the advantage of W_2 over L^2 and \mathcal{H}^{-1} in terms of its convexity property: while W_2^2 is globally convex, L^2 and \mathcal{H}^{-1} are only convex locally.

In fact, the case of $\phi(\mathbf{x}) = (2\pi)^{-d/2}e^{-\frac{1}{2}|\mathbf{x}|^2}$ is well-known in the statistics literature [9]. In this case, f and g are Gaussian densities with mean-covariance $(\lambda\bar{\mathbf{x}}, \Sigma)$ and $(\lambda_g\bar{\mathbf{x}}_g, \Sigma_g)$

respectively. The W_2^2 distance between them is simplified to

$$W_2^2(f, g) = |\lambda\bar{\mathbf{x}} - \lambda_g\bar{\mathbf{x}}_g|^2 + \text{Tr}(\Sigma + \Sigma_g - 2(\Sigma^{1/2}\Sigma_g\Sigma^{1/2})^{1/2}). \quad (35)$$

This shows that $W_2^2(f, g)$ is convex with respect to the difference of the mean and variance of the two Gaussian densities. This fact makes the quadratic Wasserstein metric extremely useful for inverse matching of Gaussian densities.

5.2 Inverse transport with W_2

The simple calculations we just had can turn out to be very useful in solving some inverse matching problems.

Transport in homogeneous flow. Let us consider the transport of a physical quantity ϕ in a given uniform flow \mathbf{v} . The evolution of the quantity is modeled by the following transport equation:

$$\frac{\partial\psi}{\partial t} + \mathbf{v} \cdot \nabla\psi = 0, \quad \text{in } \mathbb{R}_+ \times \mathbb{R}^d, \quad \psi(0, \mathbf{x}) = \phi(\mathbf{x}), \quad \text{in } \mathbb{R}^d. \quad (36)$$

It is straightforward to check that the solution to this transport equation at time $t = \lambda$ is given as

$$\psi(\lambda, \mathbf{x}) = \phi(\mathbf{x} - \lambda\mathbf{v}). \quad (37)$$

For a given function ϕ , we are interested in finding from the datum $g := \phi(\mathbf{x} - \lambda_g\mathbf{v}_g)$ the unknown flow \mathbf{v} and the travel distance λ by matching the predicted datum f with g under the W_2 metric. The result in (31) then shows that $W_2^2(f, g)$ is convex with respect to $\lambda\mathbf{v}$. More precisely,

$$W_2^2(f, g) = |\lambda\mathbf{v} - \lambda_g\mathbf{v}_g|^2. \quad (38)$$

Therefore the inverse matching problem of determining $\lambda\mathbf{v}$ from given data is a convex problem under the W_2 metric.

Note that since the dependence of $W_2(f, g)$ on λ and \mathbf{v} is only through the product, we can generalize our nonlinear model (37) by making the flow more complicated. One example of such generalization is to make the change

$$\lambda\mathbf{v} \rightarrow \sum_{j=1}^J \lambda_j\mathbf{v}_j,$$

for any J given flow $\{\mathbf{v}_j\}_{j=1}^J$. In this case, $W_2^2(f, g)$ is convex with respect to $\lambda_j\mathbf{v}_j - \lambda_{g_j}\mathbf{v}_{g_j}$ ($1 \leq j \leq J$).

Reconstruction from projections. The convexity of W_2^2 with respect to translation of signals is also useful in recovering locations of objects in high-dimensional space from (possibly random and noisy) projections. Let $P_j\phi$ be the j -th projection of ϕ on a collection of K_j coordinates $(x_{k_1}, \dots, x_{k_{K_j}}) \in \mathbb{R}^{K_j}$, that is,

$$(P_j\phi)(x_{k_1}, \dots, x_{k_{K_j}}) := \int_{\mathbb{R}^{d-K_j}} \phi(\mathbf{x}) dx_{k_{K_j}+1} \cdots dx_{k_d}.$$

We assume that data from J such projections are collected, and that each coordinate in \mathbb{R}^d has been projected onto at least once in the J projections (since otherwise the reconstruction will be nonunique). Let f_j and g_j be the j -th projections of $\phi(\mathbf{x} - \bar{\mathbf{x}})$ and $\phi(\mathbf{x} - \bar{\mathbf{x}}_g)$ respectively. Then from (31), we see that the functional

$$\Phi(\bar{\mathbf{x}}, \bar{\mathbf{x}}_g) = \sum_{j=1}^J W_2^2(f_j, g_j) = \sum_{j=1}^J \Pi(j) (\bar{x}_j - \bar{x}_{gj})^2$$

with $\Pi(j)$ the number of times that coordinate x_j is included in the J projections ($1 \leq \Pi(j) \leq J$), is globally convex with respect to each $\bar{x}_j - \bar{x}_{gj}$. Therefore, locating an objects from measured projections is a convex problem in the W_2 framework.

5.3 W_2 -based deconvolution of localized sources

Let us consider the linear matching problem (10) where A is a linear convolution operator defined as:

$$f(m(\mathbf{x})) \equiv Am(\mathbf{x}) := \int_{\mathbb{R}^d} K(\mathbf{x} - \mathbf{y})m(\mathbf{y})d\mathbf{y}. \quad (39)$$

This type of operators appear in many areas of applications, such as signal and image processing and optical imaging, where K serves as the model of the point spread function of some given physical systems [5, 10, 17, 27, 32]. The measured signal f can be viewed as output of the system for the input source m .

In many applications, A is highly smoothing, meaning that its singular values decay very fast. Inverting A to reconstruct all the information in m is often impossible due to the noise presented in the data. Here we are interested in reconstructing highly localized sources, that is, sources that have their total mass supported in a small part of their supports. Point sources are such sources and are of great importance in many practical applications. In general, let $0 < \varepsilon \ll 1$ be given. We introduce

$$\mathcal{M}_{\mathbf{y}, \varepsilon} = \{m(\mathbf{x}) \geq 0 \mid \exists B_\varepsilon(\mathbf{y}) \subseteq \text{supp}(m) \text{ s.t. } \int_{B_\varepsilon(\mathbf{y})} m(\mathbf{x})d\mathbf{x} \geq (1 - \varepsilon^{d+1}) \int_{\mathbb{R}^d} m(\mathbf{x})d\mathbf{x}\}$$

where $B_\varepsilon(\mathbf{y})$ denote the ball of radius ε centered at \mathbf{y} . Functions in $\mathcal{M}_{\mathbf{y}, \varepsilon}$ have their total mass concentrated in a ball of radius ε , and are therefore highly localized.

It is straightforward to show the following result.

Theorem 5.1. *Let f and g be generated from (39) with m_f and m_g respectively. Assume that m_f and m_g have the same total mass. (i) For any kernel function $K(\mathbf{x})$ with finite total mass, if*

$$m_\zeta(\mathbf{x}) = \delta(\mathbf{x} - \bar{\mathbf{x}}_\zeta), \quad \zeta \in \{f, g\}$$

then we have that

$$W_2^2(f, g) = |\bar{\mathbf{x}}_f - \bar{\mathbf{x}}_g|^2. \quad (40)$$

(ii) For any kernel function $K(\mathbf{x}) \in \mathcal{C}^2(\mathbb{R}^d)$, if $m_\zeta \in \mathcal{C}^2(\mathbb{R}^d) \cap \mathcal{M}_{\bar{\mathbf{x}}_\zeta, \varepsilon}$, then

$$W_2^2(f, g) = |\bar{\mathbf{x}}_f - \bar{\mathbf{x}}_g|^2 + \mathcal{O}(\varepsilon^{d+1}). \quad (41)$$

Proof. (i) follows from the fact that in this setting, $\zeta(\mathbf{x}) := Am_\zeta(\mathbf{x}) = K(\mathbf{x} - \bar{\mathbf{x}}_\zeta)$, $\zeta \in \{f, g\}$, which is only a translation of K by $\bar{\mathbf{x}}_\zeta$. To prove (ii), we first observe that inside the $B_\varepsilon(\bar{\mathbf{x}}_\zeta)$ we can write

$$\begin{aligned} m(\mathbf{y}) &= m(\bar{\mathbf{x}}_\zeta) + (\mathbf{y} - \bar{\mathbf{x}}_\zeta) \cdot \nabla m(\bar{\mathbf{x}}_\zeta) + \mathcal{O}(|\mathbf{y} - \bar{\mathbf{x}}_\zeta|^2), \\ K(\mathbf{x} - \mathbf{y}) &= K(\mathbf{x} - \bar{\mathbf{x}}_\zeta) + (\bar{\mathbf{x}}_\zeta - \mathbf{y}) \cdot \nabla K(\mathbf{x} - \bar{\mathbf{x}}_\zeta) + \mathcal{O}(|\bar{\mathbf{x}}_\zeta - \mathbf{y}|^2). \end{aligned}$$

Therefore we have

$$\begin{aligned} \zeta(\mathbf{x}) &= \int_{\mathbb{R}^d} K(\mathbf{x} - \mathbf{y})m(\mathbf{y})d\mathbf{y} = \int_{B_\varepsilon(\bar{\mathbf{x}}_\zeta)} K(\mathbf{x} - \mathbf{y})m(\mathbf{y})d\mathbf{y} + \mathcal{O}(\varepsilon^{d+1}) \\ &= m(\bar{\mathbf{x}}_\zeta)K(\mathbf{x} - \bar{\mathbf{x}}_\zeta)|B_\varepsilon(\bar{\mathbf{x}}_\zeta)| + \mathcal{O}(\varepsilon^{d+1}). \end{aligned}$$

This gives that f and g are simply perturbations of the translations of K , which then leads to the desired result. \square

Deconvolution of highly localized sources. The first important consequence of this result, described in part (i) of the theorem, is that the deconvolution from datum g , with an *arbitrary* kernel K , to recover the location of a point source is a convex problem under the W_2 metric; see (40). This is a simple but not obvious observation because once the source is parameterized in terms of the location, the convolution problem becomes nonlinear (with respect to the location). Part (ii) of the theorem generalizes point source to any highly concentrated sources by saying that to the leading order, the W_2 metric allows us to have a convex objective function in recovering the location of a highly localized object.

Deconvolution from diffusive environments. One aspect of the result that is surprising is that it does not require the convolution kernel to take a specific form, for instance, to come from a physical system governed by the transport phenomenon such as what we described in the previous subsection. This allows us to study deconvolution with smoothing kernels that could serve as models for the propagation of the information in m in a diffusive environment. A kernel of particular importance is the Gaussian kernel

$$K(\mathbf{x}) = \frac{\Lambda_d}{|\Sigma_K|^{1/2}} e^{-\frac{1}{2}\mathbf{x}^t \Sigma_K^{-1} \mathbf{x}}, \quad (42)$$

with $\Sigma_K \in \mathbb{R}^{d \times d}$ symmetric and positive-definite and the constant $\Lambda_d = (2\pi)^{-d/2}$. For this specific kernel, we could make the asymptotic calculation in the theorem more precise. Let us take the function m to be of Gaussian type, as an approximation to the localized source we discussed in the theorem (when we make the covariance matrix small). More precisely, we take

$$m_\zeta(\mathbf{x}) = \frac{\Lambda_d}{|\Sigma_{m_\zeta}|^{1/2}} e^{-\frac{1}{2}(\mathbf{x} - \mathbf{x}_{m_\zeta})^t \Sigma_{m_\zeta}^{-1} (\mathbf{x} - \mathbf{x}_{m_\zeta})}, \quad \zeta \in \{f, g\}$$

with Σ_{m_f} and Σ_{m_g} symmetric and positive-definite. Then it is straightforward to verify that f and g are of the form:

$$\zeta(\mathbf{x}) = \frac{\Lambda_d}{|\Sigma_K + \Sigma_{m_\zeta}|^{1/2}} e^{-\frac{1}{2}(\mathbf{x} - \mathbf{x}_{m_\zeta})^t (\Sigma_K + \Sigma_{m_\zeta})^{-1} (\mathbf{x} - \mathbf{x}_{m_\zeta})}.$$

Following (35), we have that

$$W_2^2(f, g) = |\mathbf{x}_f - \mathbf{x}_g|^2 + \text{Tr}(\Sigma_f + \Sigma_g - 2(\Sigma_f^{1/2}\Sigma_g\Sigma_f^{1/2})^{1/2}),$$

where $\Sigma_\zeta = \Sigma_K + \Sigma_{m_\zeta}$. To mimic the case of localized source, we perform the rescaling $\Sigma_{m_f} \rightarrow \varepsilon_f^2 \Sigma_{m_f}$ and $\Sigma_{m_g} \rightarrow \varepsilon_g^2 \Sigma_{m_g}$ with ε_f and ε_g both small (but not necessarily the same). Then $W_2^2(f, g)$ simplifies to

$$W_2^2(f, g) = |\mathbf{x}_f - \mathbf{x}_g|^2 + \frac{1}{2} \text{Tr} \left(\Sigma_K (\varepsilon_f^2 \Sigma_K^{-1} \Sigma_{m_f} - \varepsilon_g^2 \Sigma_K^{-1} \Sigma_{m_g})^t (\varepsilon_f^2 \Sigma_K^{-1} \Sigma_{m_f} - \varepsilon_g^2 \Sigma_K^{-1} \Sigma_{m_g}) \right) + \frac{3}{8} \varepsilon_g^4 \text{Tr} \left(\Sigma_K (\Sigma_K^{-1} \Sigma_{m_g})^t (\Sigma_K^{-1} \Sigma_{m_g}) \right) + \mathcal{O}(\varepsilon_f^6 + \varepsilon_g^6 + \varepsilon_f^2 \varepsilon_g^4 + \varepsilon_f^4 \varepsilon_g^2). \quad (43)$$

Taking the shapes of m_f and m_g to be the same, that is, $\Sigma_{m_f} = \Sigma_{m_g}$, we see that $W_2^2(f, g)$ is convex with respect to both $\mathbf{x}_{m_f} - \mathbf{x}_{m_g}$ and $\varepsilon_f^2 - \varepsilon_g^2$ at the leading orders. This fact gives us the opportunity to deconvolve from Gaussian kernels to reconstruct stably the location and relative size of localized sources. Note, however, that due to the presence of the trace operator, $W_2^2(f, g)$ is less sensitive to the shape (or anisotropy) of the source encoded in its covariance matrix, compared to the size of the source.

6 Numerical simulations

In this section, we present some numerical simulations to demonstrate some of the main phenomena that we analyzed in this work. We consider two different inverse data matching problems.

6.1 Deconvolution under the W_2 metric

We first show some simulation results on the deconvolution problem in the one-dimensional setting. Our numerical simulations can only be done in a finite domain, so we set $\Omega = [-\ell, \ell]$. The forward operator is the convolution on Ω :

$$f(m(x)) = Am(x) := \int_{-\ell}^{\ell} K(x-y)m(y)dy, \quad (44)$$

for a kernel $K(x) > 0$ with finite integral $\int_{-2\ell}^{2\ell} K(x)dx = \langle K \rangle$. We assume that the data g is generated from a true function $m^*(x) \geq 0$ such that $\int_{-\ell}^{\ell} m^*(x)dx = \langle m^* \rangle$. This will ensure that $\int_{-\ell}^{\ell} g(x)dx = \langle K \rangle \langle m^* \rangle$. We will consider inverse matching with both noiseless and noisy data. In the case of noisy data, we add multiplicative noise to g in a way such that $\int_{-\ell}^{\ell} g_\delta(x)dx = \langle m^* \rangle$.

In Figure. 4 we show deconvolution results for the Laplacian kernel $K_L(x) = e^{-\ell|x|}$ with the L^2 , \mathcal{H}^{-1} and W_2 metrics, that is, using the objective functions $\Phi_{\mathcal{H}^0}(m)$, $\Phi_{\mathcal{H}^{-1}}(m)$ and $\Phi_{W_2}(m)$. This set of results show clearly the smoothing effect of the quadratic Wasserstein




Figure 4: Deconvolution with a Laplacian kernel with the L^2 (left), $\dot{\mathcal{H}}^{-1}$ (middle), and W_2 (right) metrics. Top row: deconvolution with noise-free data; Bottom row: deconvolution with data containing respectively 2%, 10%, and 10% random noise.

metric: in both the noise-free and noisy data cases, W_2 give smoothed out reconstruction compared to its L^2 counterpart. The order of the smoothing is very similar to that of the $\dot{\mathcal{H}}^{-1}$ norm, although the exact reconstructions are quite different. Compared to $\dot{\mathcal{H}}^{-1}$, W_2 reconstructions seem to have slightly more variations. Both the $\dot{\mathcal{H}}^{-1}$ and the W_2 reconstructions can tolerate strong high-frequency random noise, while the L^2 reconstructions break down quickly at the relatively low noise level. One can certainly add a regularization mechanism to stabilize the L^2 reconstructions here. However, our objective here is not to say that W_2 reconstruction is better than L^2 reconstruction, but to demonstrate the smoothing effect of the quadratic Wasserstein metric. Therefore we do not include any explicit regularization to the L^2 reconstruction (although our discretization of the forward problem, the adjoint problem as well as the mismatch between the forward and adjoint discretization, indeed introduce a regularization effect).

We have done extensive numerical experiments on this problem with data at different noise levels and different convolution kernels. The same type of smoothing effect are observed in all the numerical tests with the quadratic Wasserstein metric. Interested readers are refer to [11] for more detailed description of such results. As an example, we show in Figure 5 deconvolution results from the kernel $K_I(x) = \frac{1}{1+|x|}$.

6.2 Inverse matching for an elliptic PDE problem

The second set of numerical simulations focuses on an inverse coefficient problem for the diffusion equation. Let $\Omega \subset \mathbb{R}^d$ be a bounded domain with boundary $\partial\Omega$ and consider the




Figure 5: Deconvolution with the kernel $K_I(x) = \frac{1}{1+|x|}$ with the L^2 (left), \mathcal{H}^{-1} (middle), and W_2 (right) metrics. Top row: deconvolution with noise-free data; Bottom row: deconvolution with data containing respectively 2%, 10%, and 10% random noise.

following second-order elliptic partial differential equation with Robin boundary condition:

$$-\nabla \cdot \gamma u(\mathbf{x}) + \sigma(\mathbf{x})u(\mathbf{x}) = 0, \quad \text{in } \Omega, \quad \mathbf{n} \cdot \gamma \nabla u(\mathbf{x}) + \kappa u = h(\mathbf{x}), \quad \text{on } \partial\Omega. \quad (45)$$

Physically this equation can be used to describe the diffuse of particles, generated from source $h(\mathbf{x}) \geq 0$, in an absorbing environment, such as propagation of near infra-red photons in biological tissues, with $\gamma(\mathbf{x}) > 0$ and $\sigma(\mathbf{x}) > 0$ respectively the diffusion and absorption coefficients of the environment.

For applications in quantitative photoacoustic imaging [4, 29], we are interested in reconstructing the absorption coefficient σ in the equation (45) from data that we measure in an experiment of the following form:

$$f(\sigma) := \sigma(\mathbf{x})u(\mathbf{x}), \quad \mathbf{x} \in \Omega \quad (46)$$

where $u(\mathbf{x})$ is the solution to the diffusion equation. Due to the fact that the diffusion solution u depends on the unknown σ in a nonlinear way, this inverse problem is nonlinear.

Let us assume that Ω is smooth, so that the diffusion equation (45) admits a unique solution that is sufficiently regular when γ , σ and h are sufficiently regular. Moreover, from standard theory of elliptic partial differential equations [16], we conclude that u is non-negative and bounded from above when the boundary source $h(\mathbf{x})$ is so. Therefore $f(\sigma)$ is nonnegative and bounded from above. We therefore can treat $f(\sigma)$ as a probability density for a given σ .

We first show in Figure 6 some typical inversion results for the problem in the one-dimensional case where $\Omega = (0, 1)$. For simplicity, we set $\gamma = 0.02$ in all the simulations




Figure 6: Inversion for σ in $\Omega = (0, 1)$ with noise-free (top row) and noisy data (bottom row) under the L^2 (left), \mathcal{H}^{-1} (middle) and W_2 (right) metrics. The noise level in the bottom row is 12% for each case.

we present here. Our extensive numerical tests show that the exact value of γ has only a negligible impact on the inversion results. We again observe the smoothing effect of the W_2 metric as in the deconvolution example, and this smoothing effect makes the inversion very robust against high-frequency random noise in the data. This smoothing effect is independent of the dimension of the spatial domain, as seen from the analysis in the previous sections of the paper and observed in two-dimensional inversion results shown in Figure 7.

7 Concluding remarks

In this work, we performed analytical and computational studies on the effect of the quadratic Wasserstein metric in solving linear and nonlinear inverse data matching problems. We demonstrate, through analysis in the asymptotic regime, that the quadratic Wasserstein metric has a smoothing effect on the inversion results, that is, at a given noise level, matching results with W_2 is smoother than those with the L^2 metric. One can see this smoothing effect at the true inverse solution or at a local minimizer of the objective function (which could be far away from the true matching solution) when a numerical minimization approach is used to solve the matching problem. The order of the smoothing effect is the same as that of the \mathcal{H}^{-1} metric.

The smoothing effect of the quadratic Wasserstein metric in the non-asymptotic regime indicates that the landscape of the Wasserstein objective function and that of the L^2 objective function are quite different. In other words, if we solve the inverse matching problems




Figure 7: Inversion for σ in two-dimensional case in the domain $\Omega = (0, 1) \times (0, 1)$ with data containing 10% random noise. Shown from left to right are the true coefficient, the reconstructions with the L^2 , \mathcal{H}^{-1} and W_2 metrics respectively.

with numerical minimization, the trajectory of the minimizing sequence based on the L^2 metric and that based on the Wasserstein metric is different. Even if the two minimization algorithms start at the same initial guess and converge to the same minimizer, their paths between the starting point and the final point are different. Finding good ways to precisely characterize the differences between the landscapes of the two minimization problems will be the subject of future works.

Even though the smoothing effect of W_2 is very similar to that of the \mathcal{H}^{-1} metric, we showed, through analyzing some simple finite-dimensional inverse matching problems, such as the deconvolution of point sources from given kernels, that W_2 has better convexity than \mathcal{H}^{-1} and the L^2 norm in the non-asymptotic regime. Characterizing such advantages in more practically useful situations would be an interesting project.

The quadratic Wasserstein metric we studied in this paper requires that the two data to be compared, f and g , be both probability measures (and therefore be nonnegative) and they have the same total mass. This requirement severely limits the applicability of the Wasserstein metric to more general data matching problems. To use the Wasserstein metrics for data that are not nonnegative, for instance in data matching for seismic imaging applications [7, 13, 14, 25, 26, 35, 36], one can perform some rescaling on the data to make them nonnegative. It seems, at least on the computational level, that the main features of the Wasserstein metrics are preserved by reasonable rescaling strategies such as $f \mapsto \alpha + \beta f$ and $f \mapsto e^{\beta f}$ (α and β being appropriately chosen constants). There have been also attempts to generalize the original optimal transport framework to deal with signed data [2, 22], although these have remained mostly on the theoretical level so far, except in [24] where a strategy for dealing with non-positive data is proposed and applied to the seismic imaging problem. To compare two nonnegative signals that do not have the same total mass, one can either use a normalization procedure, for instance by dividing the signals with their respective total mass so that they both have total mass 1 (c.g. [7, 14, 26, 36]), or consider generalized Wasserstein distances induced by unbalanced optimal transport. In terms of the later, the Wasserstein-Fisher-Rao metric [8, 23] has been studied extensively in the literature. Analyzing the behavior of the quadratic Wasserstein metric in the aforementioned more

complicated setups is an interesting topic for future research.

Acknowledgments

We would like to thank the anonymous referees for providing useful comments on the first draft of this work and pointing out several related references. Their inputs helped us improve the quality of this work. This work is partially supported by the National Science Foundation through grants DMS-1620396, DMS-1620473, DMS-1913129, and DMS-1913309.

References

- [1] L. AMBROSIO AND N. GIGLI, *A user guide to optimal transport*, in *Modelling and Optimisation of Flows on Networks*, Springer, 2013, pp. 1–155.
- [2] L. AMBROSIO, E. MAININI, AND S. SERFATY, *Gradient flow of the Chapman Rubinstein Schatzman model for signed vortices*, *Ann. Inst. H. Poincaré C*, 28 (2011), pp. 217–246.
- [3] G. BAL AND K. REN, *Physics-based models for measurement correlations. application to an inverse Sturm-Liouville problem*, *Inverse Problems*, 25 (2009). 055006.
- [4] ———, *Multi-source quantitative PAT in diffusive regime*, *Inverse Problems*, 27 (2011). 075003.
- [5] W. E. BLASS AND G. W. HALSEY, *Deconvolution of Absorption Spectra*, Academic Press, 1981.
- [6] R. BROSSIER, S. OPERTO, AND J. VIRIEUX, *Which data residual norm for robust elastic frequency-domain full waveform inversion*, *Geophysics*, 75 (2010), pp. R37–R46.
- [7] J. CHEN, Y. CHEN, H. WU, AND D. YANG, *The quadratic Wasserstein metric for earthquake location*, *J. Comput. Phys.*, 373 (2018), pp. 188–209.
- [8] L. CHIZAT, G. PEYRÉ, B. SCHMITZER, AND F. X. VIALARD, *An interpolating distance between optimal transport and Fisher-Rao metrics*, *Found. Comput. Math.*, 18 (2016), pp. 1–44.
- [9] J. DELON AND A. DESOLNEUX, *A wasserstein-type distance in the space of gaussian mixture models*, arXiv:1907.05254v1, (2019).
- [10] V. DIMRI, *Deconvolution and Inverse Theory : Application to Geophysical Problems*, *Methods in Geochemistry and Geophysics*, Elsevier, Amsterdam, 1992.
- [11] W. DING, Q. DU, AND K. REN, *Optimal transport for computational inverse problems*, Preprint, (2019).
- [12] H. W. ENGL, M. HANKE, AND A. NEUBAUER, *Regularization of Inverse Problems*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1996.
- [13] B. ENGQUIST AND B. D. FROESE, *Application of the Wasserstein metric to seismic signals*, *Commun. Math. Sci.*, 12 (2014), pp. 979–988.

- [14] B. ENGQUIST, B. D. FROESE, AND Y. YANG, *Optimal transport for seismic full waveform inversion*, Commun. Math. Sci., 14 (2016), pp. 2309–2330.
- [15] B. ENGQUIST AND Y. YANG, *Seismic inversion and the data normalization for optimal transport*, Methods and Applications of Analysis, (2019).
- [16] D. GILBARG AND N. S. TRUDINGER, *Elliptic Partial Differential Equations of Second Order*, Springer-Verlag, Berlin, 2000.
- [17] R. C. GONZALEZ AND R. E. WOODS, *Digital Image Processing*, Prentice Hall, 2002.
- [18] S. KOLOURI, S. R. PARK, M. THORPE, D. SLEPCEV, AND G. K. ROHDE, *Optimal mass transport: signal processing and machine-learning applications*, IEEE Signal Processing Magazine, 34 (2017), pp. 43–59.
- [19] J. LELLMANN, D. LORENZ, C. SCHÖNLIEB, AND T. VALKONEN, *Imaging with Kantorovich-Rubinstein discrepancy*, SIAM J. Imag. Sci., 7 (2014), pp. 2833–2859.
- [20] L. LI, A. VIDARD, F.-X. LE DIMET, AND J. MA, *Topological data assimilation using Wasserstein distance*, Inverse Problems, 35 (2019). 015006.
- [21] G. LOEPER, *Uniqueness of the solution to the Vlasov-Poisson system with bounded density*, J. Math. Pures Appl., 86 (2006), pp. 68–79.
- [22] E. MAININI, *A description of transport cost for signed measures*, J. Math. Sci., 181 (2012), pp. 837–855.
- [23] L. MÉTIVIER, A. ALLAIN, R. BROSSIER, Q. MÉRIGOT, E. OUDET, AND J. VIRIEUX, *Optimal transport for mitigating cycle skipping in full waveform inversion: a graph space transform approach*, Geophysics, 83 (2018), pp. R515–R540.
- [24] L. MÉTIVIER, R. BROSSIER, Q. MÉRIGOT, AND E. OUDET, *A graph space optimal transport distance as a generalization of L^p distances: application to a seismic imaging inverse problem*, Inverse Problems, 35 (2019). 085001.
- [25] L. MÉTIVIER, R. BROSSIER, Q. MÉRIGOT, E. OUDET, AND J. VIRIEUX, *Measuring the misfit between seismograms using an optimal transport distance: application to full waveform inversion*, Geophys. J. Int., 205 (2016), pp. 345–377.
- [26] ———, *An optimal transport approach for seismic tomography: application to 3D full waveform inversion*, Inverse Probl., 32 (2016). 115008.
- [27] M. J. NASSE AND J. C. WOEHLE, *Realistic modeling of the illumination point spread function in confocal scanning optical microscopy*, J. Opt. Soc. Am. A., 27 (2010), pp. 295–302.
- [28] R. PEYRE, *Comparison between W_2 distance and \dot{H}^{-1} norm, and localization of wasserstein distance*, ESAIM: COCV, (2018).
- [29] K. REN, H. GAO, AND H. ZHAO, *A hybrid reconstruction method for quantitative photoacoustic imaging*, SIAM J. Imag. Sci., 6 (2013), pp. 32–55.
- [30] F. SANTAMBROGIO, *Optimal Transport for Applied Mathematicians: Calculus of Variations, PDEs, and Modeling*, Birkhäuser, 2015.

- [31] J. SOLOMON, R. M. RUSTAMOV, L. GUIBAS, AND A. BUTSCHER, *Wasserstein propagation for semi-supervised learning*, in Proceedings of the 31st International Conference on Machine Learning, E. P. Xing and T. Jebara, eds., 2014, pp. 306–314.
- [32] J. N. TURNER, S. LASEK, AND D. H. SZAROWSKI, *Confocal optical microscopy*, in Encyclopedia of Materials : Science and Technology, 2001, pp. 1504–1509.
- [33] C. VILLANI, *Topics in Optimal Transportation*, American Mathematical Society, Providence, RI., 2003.
- [34] C. R. VOGEL, *Computational Methods for Inverse Problems*, Frontiers in Applied Mathematics, SIAM, Philadelphia, 2002.
- [35] Y. YANG AND B. ENGQUIST, *Analysis of optimal transport and related misfit functions in full-waveform inversion*, Geophysics, 83 (2018), pp. A7–A12.
- [36] Y. YANG, B. ENGQUIST, J. SUN, AND B. D. FROESE, *Application of optimal transport and the quadratic Wasserstein metric to full waveform inversion*, Geophysics, 83 (2018), pp. R43–R62.