Wasserstein GANs with Gradient Penalty Compute Congested Transport

Tristan Milne   TMILNE@MATH.TORONTO.EDU
Department of Mathematics, University of Toronto, 40 St. George Street, Toronto Ontario Canada

Adrian Nachman   NACHMAN@MATH.TORONTO.EDU
Department of Mathematics, University of Toronto, 40 St. George Street, Toronto Ontario Canada
The Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, 10 King's College Road, Toronto Ontario Canada

Editors: Po-Ling Loh and Maxim Raginsky

Abstract

Wasserstein GANs with Gradient Penalty (WGAN-GP) are a very popular method for training generative models to produce high quality synthetic data. While WGAN-GP were initially developed to calculate the Wasserstein 1 distance between generated and real data, recent works (e.g. [23]) have provided empirical evidence that this does not occur, and have argued that WGAN-GP perform well not in spite of this issue, but because of it. In this paper we show for the first time that WGAN-GP compute the minimum of a different optimal transport problem, the so-called congested transport [7]. Congested transport determines the cost of moving one distribution to another under a transport model that penalizes congestion. For WGAN-GP, we find that the congestion penalty has a spatially varying component determined by the sampling strategy used in [12] which acts like a local speed limit, making congestion cost less in some regions than others. This aspect of the congested transport problem is new, in that the congestion penalty turns out to be unbounded and depends on the distributions to be transported, and so we provide the necessary mathematical proofs for this setting. One facet of our discovery is a formula connecting the gradient of solutions to the optimization problem in WGAN-GP to the time averaged momentum of the optimal mass flow. This is in contrast to the gradient of Kantorovich potentials for the Wasserstein 1 distance, which is just the normalized direction of flow. Based on this and other considerations, we speculate on how our results explain the observed performance of WGAN-GP. Beyond applications to GANs, our theorems also point to the possibility of approximately solving large scale congested transport problems using neural network techniques.

1. Introduction

Wasserstein GANs (WGANs) were first proposed in [2] as a means of training generative models using the Wasserstein 1 distance to measure the dissimilarity of the real and generated distributions. Recall that the Wasserstein 1 distance (or Earth Mover's Distance) between two probability distributions μ and ν on a subset Ω of $\mathbb{R}^d$ can be calculated via duality as

$$W_1(\mu, \nu) = \sup_{u \in 1\text{-Lip}(\Omega)} \mathbb{E}_{x \sim \mu}[u(x)] - \mathbb{E}_{y \sim \nu}[u(y)], \tag{1}$$

where $1\text{-Lip}(\Omega)$ is the set of 1-Lipschitz real valued functions on Ω. The authors of [2] showed that $W_1(\mu, \nu)$ has better theoretical properties than the Jenson-Shannon divergence used in the original GAN paper [11], and proposed to solve (1) by parametrizing u as a neural network u_w with parameters w. This network, dubbed the "critic", replaced the discriminator of the original GAN. However,

designing critics u_w which are 1-Lipschitz and sufficiently expressive is a non-trivial task which has only more recently seen some progress [1].

As an initial resolution to this issue, the authors of [2] used weight-clipping, but the performance of WGANs in terms of the visual quality of generated images and training stability was greatly improved in [12], wherein weight clipping was discarded in favour of solving the following optimization problem,

$$\sup_u \mathbb{E}_{x \sim \mu}[u(x)] - \mathbb{E}_{y \sim \nu}[u(y)] - \lambda \mathbb{E}_{z \sim \sigma}[(|\nabla u(z)| - 1)^2]. \tag{2}$$

Here the hard constraint in (1) that u be 1-Lipschitz (i.e. $|\nabla u| \leq 1$ almost everywhere) is replaced with a penalty term which penalizes $|\nabla u|$ being different from 1. The probability distribution σ is defined by the following procedure: independently sample $x \sim \mu$, $y \sim \nu$, and t from the uniform distribution on $[0, 1]$ to obtain the point $z \sim \sigma$ given by the formula

$$z = (1 - t)x + ty.$$

It was observed in [12] that the removal of the hard constraint enforced by weight-clipping allowed for a much more expressive class of functions at the cost of sacrificing guarantees of the functions being 1-Lipschitz. Generative models trained with this approach are called WGAN-GP (GP for gradient penalty), and they have emerged as a popular method for stably training generative models to produce high-quality synthetic images [13; 16].

The authors of [12] also suggested modifying (2) to

$$\sup_u \mathbb{E}_{x \sim \mu}[u(x)] - \mathbb{E}_{y \sim \nu}[u(y)] - \lambda \mathbb{E}_{z \sim \sigma}[(|\nabla u(z)| - 1)_+^2], \tag{3}$$

which uses a one-sided gradient penalty ($a_+ := \max(a, 0)$ for $a \in \mathbb{R}$). This is the version of WGAN-GP studied in this paper. This is because it is a more natural penalty term than the two-sided penalty for encouraging functions to be 1-Lipschitz, since it only penalizes gradients larger than one in norm. Further, it has the important advantage of being convex in u. Finally, studies (e.g. [12; 20]) have shown that it obtains equal or better performance than the two-sided penalty in (2), and it has seen repeated use (see, for example, [17; 19; 20]).

While WGAN-GP have enjoyed spectacular success, the question of whether they are actually computing the Wasserstein 1 distance has only been studied more recently in, for example, [18], [21], and [23]. In particular we were intrigued by [23], which offers empirical evidence that WGAN-GP do *not* compute $W_1(\mu, \nu)$, and, due to some issues with the Wasserstein 1 distance, argues that this might be the reason for their success. In this paper we analyse this further by establishing for the first time the quantity that WGAN-GP with one-sided penalty *do* compute: it is a congested transport distance.

We will explain precisely what this means in Section 3, but vaguely, congested transport is a branch of optimal transport theory that seeks to model the optimal flow of mass under the effects of congestion. That is, like moving through a busy city, the concentration of mass in a region affects the minimal time required to move through that region. This is distinguished from the typical optimal transport model, where the cost of moving mass from some point x to another y does not take into account the traffic along the way. The theory of congested transport for the continuous case (as opposed to the discrete one) was developed in [6] and [7]; see also Chapter 4 of [22].

The contributions of our paper are as follows:

i. Under mild assumptions on μ and ν, we establish that the optimal value (3) of WGAN-GP is equal to the minimal cost of moving μ to ν as determined by a congested transport model.

ii. We establish that the value of this minimal cost is not equal to $W_1(\mu, \nu)$ if $\mu \neq \nu$. More precisely, it is strictly larger for all $\lambda > 0$, and at best it converges to $W_1(\mu, \nu)$ like λ^{-1} as $\lambda \to \infty$, if it converges at all.

iii. Under slightly stronger assumptions on μ and ν, we prove the existence of a solution u_0 to (3) in an appropriate function space.

iv. Further, we show that there is a formula relating $\nabla u_0(x)$ to the time averaged momentum (i.e. mass times velocity) at x of the optimal mass flow for the congested transport problem. This is in contrast to the standard Wasserstein 1 framework, where the gradient of a solution to (1) is just the direction of optimal mass transport, and does not encode the speed of that transport or the amount of mass transported.

v. Finally, we show that the statement in iv also holds approximately for approximate solutions of (3), which is significant in practice since numerical algorithms for solving (3) will only produce approximate solutions.

The plan for this paper is as follows. In Section 2 we provide the basic background on congested transport needed to state our main theorems. These are given in Section 3. We discuss related work in Section 4. Proof sketches for our main results are presented in Section 5, while the details are deferred to Appendix A. We provide in Section 6.1 some intuition for how our results may explain the observed performance of WGAN-GP, and summarize the paper in Section 6.2.

2. Background on congested transport

In this section we will highlight the aspects of congested transport required to present our results. For more background see [6] and [7], or Chapter 4 of [22].

The standard optimal transport problem for probability distributions μ and ν on a subset Ω of Euclidean space $\mathbb{R}^d$ with cost c is given by

$$\inf\{\int_{\Omega \times \Omega} c(x, y)d\gamma \mid \gamma \in \mathcal{P}(\Omega \times \Omega), (\pi_x)_\# \gamma = \mu, (\pi_y)_\# \gamma = \nu\}. \tag{4}$$

Here $\mathcal{P}(\Omega \times \Omega)$ is the set of probability distributions on $\Omega \times \Omega$, so that γ can be thought of as a joint probability distribution of random variables X and Y taking values in Ω. The maps $\pi_x, \pi_y : \Omega \times \Omega \to \Omega$ are the standard projections

$$\pi_x(x, y) = x, \pi_y(x, y) = y,$$

and the pushforward measures $(\pi_x)_\# \gamma$, $(\pi_y)_\# \gamma$ are the marginals of γ. As alluded to above, the cost of moving one unit of mass from x to y is $c(x, y)$, which depends only on the initial and final position and not on the path taken between those points, nor on the presence or absence of other mass on that path. Based on this insight, and on earlier work for discrete problems [24], the theory of congested transport was developed in [7] to account for possible congestion effects. In this theory, the so-called "transport plan" $\gamma \in \mathcal{P}(\Omega \times \Omega)$ of (4) is replaced with a probability distribution Q

(called a "traffic plan") on the space of absolutely continuous curves in Ω parametrized on $[0, 1]$, a space we denote by $\mathcal{C}$. It is helpful to introduce for each $t \in [0, 1]$ the "evaluation at time t" map $e_t : \mathcal{C} \to \Omega$, which sends a curve $\omega \in \mathcal{C}$ to its position at time t, $\omega(t)$. The curve of measures $(e_t)_{\#}Q$ is then a flow of mass in time, and compatibility of Q with the source and target distributions μ and ν is enforced by requiring

$$(e_0)_{\#}Q = \mu, (e_1)_{\#}Q = \nu, \tag{5}$$

which means that the flow $(e_t)_{\#}Q$ starts at μ and ends at ν. The set of traffic plans satisfying (5) will be denoted by $\mathcal{Q}(\mu, \nu)$.

A result in [7] helps to gain intuition on such traffic plans. Define $\mathcal{C}^{x,y}$ as the subset of $\mathcal{C}$ consisting of curves that start at x and end at y. It is shown in [7] that for any $Q \in \mathcal{Q}(\mu, \nu)$ there is a transport plan γ admissible in (4) and for each $x, y \in \Omega$ a distribution $Q^{x,y}$ on $\mathcal{C}^{x,y}$ such that Q decomposes as $dQ = dQ^{x,y}d\gamma$. More precisely, for any continuous test function ϕ mapping $\mathcal{C}$ to $\mathbb{R}$,

$$\int_{\mathcal{C}} \phi(\omega)dQ = \int_{\Omega \times \Omega} \int_{\mathcal{C}^{x,y}} \phi(\omega)dQ^{x,y}d\gamma.$$

In this way, traffic plans Q not only select the initial and final positions of mass with γ, but also the paths taken between x and y with $Q^{x,y}$.

To each traffic plan Q one associates a scalar measure i_Q, called the traffic intensity, and a vector measure $\mathbf{w}_Q$, called the traffic flow. For $\phi \in C(\Omega)$ and $\xi \in C(\Omega; \mathbb{R}^d)$ scalar and vector test functions, respectively, these measures are defined by the equalities

$$\int_{\Omega} \phi di_Q = \int_{\mathcal{C}} \int_0^1 \phi(\omega(t))|\omega'(t)|dt dQ, \tag{6}$$

$$\int_{\Omega} \xi \cdot d\mathbf{w}_Q = \int_{\mathcal{C}} \int_0^1 \xi(\omega(t)) \cdot \omega'(t)dt dQ. \tag{7}$$

Heuristically, for a measurable set $E \subset \Omega$, $i_Q(E)$ represents the total mass passing through E according to Q, weighted by the length of each curve in E; in this sense i_Q is a measure of congestion. The value $\mathbf{w}_Q(E)$ also has a physical interpretation. After dividing by the total elapsed time, we can think of it as the time averaged momentum (since it has units of velocity multiplied by mass) of curves, according to Q, which pass through E.

We can now define the important notion of the cost of congested transport. To each traffic plan Q that has a traffic intensity i_Q with a density with respect to Lebesgue measure (which we will denote by $i_Q(x)$), one associates a cost given by the formula,

$$\int_{\Omega} H(x, i_Q(x))dx,$$

where $H : \Omega \times \mathbb{R} \to \mathbb{R}$ is the cost function. In general, H has a non-trivial dependency on the spatial variable x (and in our case this is important) but the example often studied (e.g. in [6]) is

$$H(x, z) = H(z) = \frac{1}{2\lambda}z^2 + |z|. \tag{8}$$

With this choice of cost, the incremental cost at congestion level $i_Q(x)$ is $H'(i_Q(x)) = \frac{1}{\lambda}i_Q(x) + 1$. As such, when congestion (i.e. $i_Q(x)$) is large, the incremental cost of adding more mass at x is

very high. Conversely, when there is no traffic (i.e. $i_Q(x) = 0$), there is still a non-zero incremental cost; this is often phrased as "cars cannot travel at infinite speeds on empty roads".

The standard congested transport problem is then to minimize the cost given by H among all traffic plans $Q \in \mathcal{Q}(\mu, \nu)$ with i_Q absolutely continuous with respect to Lebesgue measure; that is

$$\inf\{\int_\Omega H(x, i_Q(x))dx \mid Q \in \mathcal{Q}(\mu, \nu), i_Q \ll \mathcal{L}_d\}.$$

See [7] for proofs of existence of solutions to this problem, and [6] for the relationship between solutions and minimal flows.

3. Main results

3.1. Definition of the optimization problems

Here we will define more precisely the optimization problems involved in WGAN-GP. Throughout we will assume that μ and ν have densities $f(x)$ and $g(x)$ with respect to Lebesgue measure; we will see later in Section 5.1 that this implies σ also has a density, which we will denote by $\sigma(x)$.

We need to specify the space over which we are maximizing in (3). The simplest choice is $H^1(\Omega)$, the Sobolev space of functions $u : \Omega \to \mathbb{R}$ with u and its weak derivative ∇u satisfying

$$\int_\Omega (u^2(x) + |\nabla u|^2(x))dx < \infty.$$

With this set of admissible functions, the problem, which we denote as (GP_λ), is as follows:

$$\sup\{\langle u, f - g \rangle - \frac{\lambda}{2} \int_\Omega (|\nabla u| - 1)_+^2 \sigma(x)dx \mid u \in H^1(\Omega)\}, \qquad (GP_\lambda)$$

where $\langle u, f - g \rangle$ denotes the $L^2(\Omega)$ inner product. The functional in (GP_λ) is exactly that of (3) up to a rescaling of λ. We emphasize that in (GP_λ), as in later problems, we work specifically with the measure σ that corresponds precisely to the sampling scheme in [12]. The value of the supremum will be denoted as $\sup(GP_\lambda)$, and throughout this paper we will use max (or min, as appropriate) rather than sup (or inf) when the optimization problem has a solution. We use the space $H^1(\Omega)$ as it is the simplest space over which we can guarantee that the functional is finite if, say, $f, g \in L^2(\Omega)$ and $\sigma \in L^\infty(\Omega)$.

A slightly more complicated but also more natural space, in view of the penalty term, is the weighted Sobolev space $H^1(\Omega, \sigma)$. This is the space of functions u with weak derivatives ∇u having finite norm according to the weight σ, i.e.

$$\|u\|_{H^1(\Omega,\sigma)} := \left(\int_\Omega (u^2(x) + |\nabla u|^2(x)) \sigma(x)dx \right)^{1/2} < \infty.$$

As we will see later we have sharper results for the case of optimizing (3) over this space. *A priori* it is not clear that the measure σ from [12] is non-degenerate almost everywhere in Ω and so it may not serve as a reasonable weight, but in Section 5.1 we will provide conditions on f and g which guarantee this. Proceeding under the assumption that σ is a reasonable weight, we write our second problem, denoted $(\widetilde{GP}_\lambda)$ as

$$\sup\{\langle u, f - g \rangle - \frac{\lambda}{2} \int_\Omega (|\nabla u| - 1)_+^2 \sigma(x)dx \mid u \in H^1(\Omega, \sigma)\}. \qquad (\widetilde{GP}_\lambda)$$

We will show in Section 5.3 that $\langle u, f - g \rangle$ has meaning for $u \in H^1(\Omega, \sigma)$, even though u may not be in $L^2(\Omega)$. This will rely on the dependence of the sampling measure σ from [12] on the given distributions μ and ν. Aside from this interesting feature of the problem, the only difference between (GP_λ) and $(\widetilde{GP_\lambda})$ is the space over which we optimize; $H^1(\Omega, \sigma)$ is a natural choice given the structure of the gradient penalty term. However, for σ to be a reasonable weight we require some additional assumptions on f and g, so we include results for (GP_λ), valid even when these assumptions fail. We will denote the value of $(\widetilde{GP_\lambda})$ as $\sup(\widetilde{GP_\lambda})$. For background on weighted Sobolev spaces, see for instance [14] or [15].

We are now ready to state our congested transport problem, denoted (CP_λ):

$$\inf\{\int_\Omega H(x, i_Q(x))dx \mid Q \in \mathcal{Q}(\mu, \nu), i_Q \ll \mathcal{L}_d, i_Q \in L^2(\Omega)\}, \qquad (CP_\lambda)$$

where $H : \Omega \times \mathbb{R} \to \mathbb{R}$ is given by

$$H(x, z) = \begin{cases} \frac{1}{2\lambda\sigma(x)} z^2 + |z| & \sigma(x) > 0, \\ 0 & \sigma(x) = 0, z = 0, \\ +\infty & \sigma(x) = 0, z \neq 0. \end{cases} \qquad (9)$$

Remarkably, the cost H that appears in our congested transport problem is very close to the one given in (8), which is the prototypical example in the congested transport literature. The only difference is that our H contains a spatially varying component depending on σ, which in turn depends on μ and ν. One can think of σ as a local speed limit, in the sense that where σ is large the cost for a certain amount of mass flow is less, and vice versa.

3.2. Statement of the main results

Throughout we will assume that $\Omega \subset \mathbb{R}^d$ is an open, bounded, convex set with a Lipschitz boundary. The following theorem is a formal statement of our contributions i, ii, and v as listed in Section 1.

Theorem A *Suppose that f and g are probability density functions in $L^\infty(\Omega)$. Then*

1. *$\sup(GP_\lambda) = \inf(CP_\lambda) < +\infty$ (i.e. WGAN-GP compute a congested transport cost).*

2. *There exists a $C > 0$, such that for all $\lambda > 0$,*

$$\sup(GP_\lambda) \geq W_1(\mu, \nu)\left(1 + \frac{C}{\lambda} W_1(\mu, \nu)\right).$$

   In particular, $\sup(GP_\lambda) > W_1(\mu, \nu)$ for all $\lambda > 0$ whenever $W_1(\mu, \nu)$ is non-zero.

3. *There is a traffic plan Q_0 that solves (CP_λ). Moreover, any two solutions Q_0, Q_1 have the same traffic flow and traffic intensity, which are related by the equations*

$$\mathbf{w}_{Q_0} = \mathbf{w}_{Q_1}, i_{Q_0} = |\mathbf{w}_{Q_0}| = |\mathbf{w}_{Q_1}| = i_{Q_1},$$

   where $|\mathbf{w}_{Q_0}|$ and $|\mathbf{w}_{Q_1}|$ are the total variation measures of the vector measures $\mathbf{w}_{Q_0}$ and $\mathbf{w}_{Q_1}$.

6

4. *If Q_0 is a solution to (CP_λ) and if $u_0 \in H^1(\Omega)$ is an approximate solution to (GP_λ) in the sense that for some $\epsilon > 0$,*

$$\langle u_0, f - g \rangle - \frac{\lambda}{2} \int_\Omega (|\nabla u_0| - 1)_+^2 \sigma dx \geq \sup(GP_\lambda) - \epsilon,$$

then the vector density $\mathbf{w}_{Q_0}$ for the traffic flow satisfies, for some constant C,

$$\left\| \mathbf{w}_{Q_0} + \lambda \sigma(|\nabla u_0| - 1)_+ \frac{\nabla u_0}{|\nabla u_0|} \right\|_{L^2(\Omega; \mathbb{R}^d)}^2 \leq C\epsilon. \tag{10}$$

With some simple additional assumptions on f and g, we are able to obtain the existence of a solution $u_0 \in H^1(\Omega, \sigma)$ to $(\widetilde{GP_\lambda})$ and make (10) an identity. Theorem **B** is a formal statement of these results, and addresses contributions iii and iv from Section 1.

Theorem B *Assume that f and g are probability density functions in $L^\infty(\Omega)$, that $g = 0$ in a neighbourhood of the boundary of Ω, and that $\inf_{x \in \Omega} f(x) > 0$. Then*

1. *All of the claims of Theorem **A** hold,*

2. *$(\widetilde{GP_\lambda})$ has a solution,*

3. *$\sup(GP_\lambda) = \max(\widetilde{GP_\lambda})$, and*

4. *If u_0 solves $(\widetilde{GP_\lambda})$ and Q_0 solves (CP_λ), then*

$$\mathbf{w}_{Q_0} = -\lambda \sigma(|\nabla u_0| - 1)_+ \frac{\nabla u_0}{|\nabla u_0|}, \quad \nabla u_0(x) = -\left(\frac{1}{\lambda \sigma(x)} + \frac{1}{|\mathbf{w}_{Q_0}(x)|} \right) \mathbf{w}_{Q_0}(x), \tag{11}$$

the latter formula holding for almost all x such that $\mathbf{w}_{Q_0}(x) \neq 0$.

4. Related Work

Since our work connects two previously disconnected fields (i.e. congested transport and generative modelling with WGANs), we will briefly review related works from both areas here.

To our knowledge, ours are the first results in the WGAN literature determining precisely the optimal value of the objective function from [12] with one-sided penalty. More generally, several papers (e.g. [18; 21; 23]) have provided empirical evidence that WGAN-GP do not compute $W_1(\mu, \nu)$, but do not offer a precise notion of what they do compute. In particular, [18] shows that for simple discrete problems WGAN-GP tend to over-estimate the Wasserstein 1 distance. A similar observation is noted in [21], who also provide evidence that the optimal value computed by WGAN-GP tends to converge to $W_1(\mu, \nu)$ only as $\lambda \to \infty$. Statement 2 of Theorem **A** provides a theoretical explanation for both of these observations.

The subject of one versus two-sided penalties for WGAN-GP is studied in depth in [20], which shows empirically that the former results in more stable training of WGANs with less dependence on the λ parameter. Despite this, it seems that the two-sided penalty is the default for many practitioners (e.g. [16]). Our results only apply directly to the one-sided penalty, but they are useful for bounding the optimal value of the two-sided penalty using the bounds provided in [20]; a more precise analysis for the two-sided penalty may be obtainable through a convexification argument.

An interesting recent work is [3], which deals with some of the issues that arise when approximating a Kantorovich potential with neural networks, though does not consider the functional used in WGAN-GP. Since our work leaves such approximation issues unexamined, we view [3] as an important and complementary perspective.

On the congested transport side, we note that our proof of the equivalence of $\sup(GP_\lambda)$ and $\inf(CP_\lambda)$ is inspired by [6] and refinements in [5]. Indeed, arguments in [6] show this equivalence for the case of $\sigma(x) = 1$ almost everywhere. However, we were unable to find the results we needed for the σ from [12] in the literature; [6] works with the case of $\sigma = 1$, and [5] includes generalizations of these results to the case of inhomogeneous (i.e. x dependent) cost $H(x, z)$. However, they assume H is bounded above for fixed z, which is not the case for WGAN-GP (see (9)). Thus one can view our results as extending some of the work in both [6] and [5] to a new class of problems where the cost function $H(x, z)$ appearing in the congested transport problem is unbounded in x; we obtain this extension by exploiting the dependence of σ on the distributions μ and ν.

Let us also note that there are variants of congested transport which are not isotropic, in the sense that the cost function depends on the direction of traffic flow (e.g. [4]). We speculate that this might lead to interesting generalizations of WGAN-GP, but leave this for future work.

5. Proof of the main theorems

In this section we will sketch the proofs of Theorems **A** and **B**; detailed proofs are provided in Appendix A. The proof of Theorem **A** is analogous to the approach from congested transport theory (i.e. [5] and [6]). We consider a dual problem given by the following Beckmann type problem:

$$\inf\{\int_\Omega H(x, |\mathbf{w}|(x))dx \mid \mathbf{w} \in L^2(\Omega; \mathbb{R}^d), \nabla \cdot \mathbf{w} = f - g\}, \qquad (BP_\lambda)$$

where H is given in (9), and as is usual for Beckmann problems the equation $\nabla \cdot \mathbf{w} = f - g$ holds in the weak sense. In other words, $\mathbf{w}$ is admissible in (BP_λ) only if, for all $u \in H^1(\Omega)$,

$$\int_\Omega -\nabla u \cdot \mathbf{w}dx = \int_\Omega u(f - g)dx.$$

Note that since we do not impose conditions on the behaviour of u at the boundary of Ω, this imposes a no-flux condition on $\mathbf{w}$. We then prove that $\sup(GP_\lambda)$ equals $\inf(BP_\lambda)$, and that the latter is equal to $\inf(CP_\lambda)$. The relationship (10) between approximate solutions of (GP_λ) and (CP_λ) is then established using properties of the Legendre dual of the function $H(x, z)$ defined in (9).

To prove Theorem **B**, we show in Proposition 3 that $\sigma(x)$ is comparable to the distance to the boundary function

$$\text{dist}(x, \partial\Omega) := \inf_{y \in \partial\Omega} |x - y|.$$

This establishes some properties of $H^1(\Omega, \sigma)$ which we need for our proof, such as a Poincaré inequality and density of smooth functions up to the boundary. The former is used to establish the existence of a minimizer to $(\widetilde{GP_\lambda})$, and the latter is used to establish $\sup(GP_\lambda) = \max(\widetilde{GP_\lambda})$. Finally, the equalities in (11) are obtained, essentially, by evaluating (10) at a solution to $(\widetilde{GP_\lambda})$.

5.1. Properties of σ

We begin by collecting several properties of the measure σ which will be needed in the proofs of Theorems **A** and **B**. If x, y, and t are sampled independently (as is done in [12]), σ is the probability distribution defined by the formula

$$\sigma = \pi_{\#}(U[0, 1] \otimes \mu \otimes \nu),$$

where $U[0, 1]$ is the uniform measure on $[0, 1]$, $U[0, 1] \otimes \mu \otimes \nu$ is the product measure, and $\pi :$ $[0, 1] \times \Omega \times \Omega \to \Omega$ is the map

$$\pi(t, x, y) = (1 - t)x + ty.$$

The following lemma guarantees that if μ and ν have densities with respect to Lebesgue measure on $\mathbb{R}^d$ then σ also has a density. We also provide a formula for this density.

Lemma 1 *If $\mu = f(x)dx$ and $\nu = g(x)dx$ then σ has a density given by the formula*

$$\sigma(z) = \int_0^1 \int_\Omega f\left(\frac{z - ty}{1 - t}\right) g(y)(1 - t)^{-d} dy dt. \tag{12}$$

The proof is a straightforward calculation; see Appendix A.1 for details.

We next establish that σ is bounded when f and g are bounded. In particular this shows that the functional in (GP_λ) is finite over $H^1(\Omega)$ for bounded f and g. See Appendix A.1 for a full proof.

Lemma 2 *If f and g are probability density functions in $L^\infty(\Omega)$, then so is σ.*

Finally we establish that under the assumptions of Theorem **B**, σ is comparable to $\mathrm{dist}(x, \partial\Omega)$.

Proposition 3 *Assume that f and g are probability density functions in $L^\infty(\Omega)$, that $g = 0$ in a neighbourhood of the boundary of Ω, and that $\inf_{x \in \Omega} f(x) > 0$. Then there exists a constant C such that for all $x \in \Omega$,*

$$\frac{1}{C} \mathrm{dist}(x, \partial\Omega) \le \sigma(x) \le C \, \mathrm{dist}(x, \partial\Omega).$$

Proposition 3 is proved in Appendix A.1 in Lemmas 11, 12 and 13. The lower bound on f is sufficient for the lower bound on σ, while convexity of Ω and having g vanish in a neighbourhood of $\partial\Omega$ is sufficient for the upper bound. Note that in this case σ vanishes on $\partial\Omega$, and hence the cost $H(x, z)$ is unbounded as x approaches the boundary.

5.2. Proof of Theorem A

For convenience, we denote the functional arising in (GP_λ) as $J : H^1(\Omega) \times L^2(\Omega; \mathbb{R}^d) \to \mathbb{R}$,

$$J(u, \mathbf{p}) = \langle u, g - f \rangle + \frac{\lambda}{2} \int_\Omega (|\mathbf{p}| - 1)_+^2 \sigma(x) dx. \tag{13}$$

We note that since $\sigma \in L^\infty(\Omega)$ (see Lemma 2), J does indeed take finite values. The problem (GP_λ) can then be written as

$$- \inf_{u \in H^1(\Omega)} J(u, \nabla u).$$

The proof of Theorem **A** rests on the following inequality, which takes into account the dependence of σ on the given probability density functions f and g. A full proof is given in Appendix A.2.

Lemma 4 *If f and g are probability density functions in $L^\infty(\Omega)$, and σ is the density given in* (12)*, then for all $u \in H^1(\Omega)$ we have*

$$|\langle u, g - f \rangle| \leq \operatorname{diam}(\Omega) \left(\int_\Omega |\nabla u|^2 \sigma(x) dx \right)^{1/2}, \tag{14}$$

where $\operatorname{diam}(\Omega) = \sup\{|x - y| \mid x, y \in \Omega\}$. Further, $\sup(GP_\lambda)$ *is finite.*

5.2.1. EQUALITY OF $\sup(GP_\lambda)$ AND $\inf(BP_\lambda)$

Here we address the equality of $\sup(GP_\lambda)$ and $\inf(BP_\lambda)$. The proof also establishes the existence of a unique solution to (BP_λ).

Proposition 5 *If f and g are probability density functions in $L^\infty(\Omega)$, then*

$$\sup(GP_\lambda) = \inf(BP_\lambda). \tag{15}$$

Furthermore, (BP_λ) has a unique minimizer.

The proof consists of verifying the conditions needed for strong duality in convex analysis (see, for example, [9]). To do so, we rely on Lemma 4, as well as convexity and continuity of J. See Appendix A.2 for the details.

5.2.2. EQUIVALENCE OF (BP_λ) AND (CP_λ)

The following result, with Proposition 5, will complete the proof of statements 1 and 3 of Theorem A. After using the existence of a solution to (BP_λ), proved in Proposition 5, the proof for unbounded $H(x, z)$ is essentially the same as in [5] and relies on the monotonicity of $H(x, z)$ in $|z|$.

Proposition 6 *If f and g are probability density functions in $L^\infty(\Omega)$, we have*

$$\min(BP_\lambda) = \inf(CP_\lambda). \tag{16}$$

Moreover, (CP_λ) has a solution, and any solution Q_0 to (CP_λ) is related to the unique solution $\mathbf{w}_0$ to (BP_λ) through the equations

$$\mathbf{w}_{Q_0} = \mathbf{w}_0, \quad i_{Q_0} = |\mathbf{w}_0|. \tag{17}$$

Proof sketch We first show that for each Q admissible in (CP_λ), the traffic flow $\mathbf{w}_Q$ is admissible for (BP_λ) and satisfies $i_Q \geq |\mathbf{w}_Q|$; this, together with monotonicity of $H(x, z)$ in z, will establish

$$\inf(CP_\lambda) \geq \inf(BP_\lambda). \tag{18}$$

To prove the reverse, we show that for every $\mathbf{w}$ admissible in (BP_λ) there is a traffic plan Q with $i_Q \leq |\mathbf{w}|$ via Theorem 4.10 of [22]. The proof of existence of a solution Q_0 and its relation to $\mathbf{w}_0$ then follows from this argument together with Proposition 5. See Appendix A.2 for a details. ∎

Having verified statements 1 and 3 of Theorem A, we now turn to statements 2 and 4.

5.2.3. INEQUALITY BETWEEN $\sup(GP_\lambda)$ AND $W_1(\mu, \nu)$

Here we will establish, using Proposition 5, that the value computed by solving (GP_λ) is strictly larger than $W_1(\mu, \nu)$ for all $\lambda > 0$ if $\mu \neq \nu$, and that at best $\sup(GP_\lambda)$ decays to $W_1(\mu, \nu)$ like λ^{-1} as $\lambda \to +\infty$. This is statement 2 of Theorem A.

Lemma 7 *If f and g are probability density functions in $L^\infty(\Omega)$, there exists a constant $C > 0$ depending on $\|\sigma\|_{L^\infty(\Omega)}$ such that*

$$\sup(GP_\lambda) \geq W_1(\mu, \nu)\left(1 + \frac{C}{\lambda}W_1(\mu, \nu)\right).$$

Proof sketch The proof follows by bounding

$$\inf(BP_\lambda) \geq \inf_{\nabla \cdot \mathbf{w} = \mu - \nu} \int_\Omega \frac{1}{2\lambda\sigma}|\mathbf{w}|^2 dx + \inf_{\nabla \cdot \mathbf{w} = \mu - \nu} \int_\Omega |\mathbf{w}|(x)dx,$$

and estimating each term individually; the second of these is known to be $W_1(\mu, \nu)$ (see, for example, [22] Chapter 4). See Appendix A.2 for full details. ∎

5.2.4. RELATIONSHIP BETWEEN APPROXIMATE SOLUTIONS OF (GP_λ) AND SOLUTIONS OF (CP_λ)

Statement 4 of Theorem A is an immediate consequence of the following proposition.

Proposition 8 *If f and g are probability density functions in $L^\infty(\Omega)$ and $\mathbf{w}_0$ is a solution to (BP_λ), then for all $u \in H^1(\Omega)$,*

$$\sup(GP_\lambda) + J(u, \nabla u) \geq \frac{1}{2\lambda\|\sigma\|_{L^\infty(\Omega)}}\left\|\mathbf{w}_0 + \lambda\sigma(|\nabla u| - 1)_+\frac{\nabla u}{|\nabla u|}\right\|_{L^2(\Omega)}^2. \tag{19}$$

See Appendix A.2 for a detailed proof, which comes from estimates on H and its Legendre dual.

5.3. Proof of Theorem B

We begin by recording some properties of $H^1(\Omega, \sigma)$, necessary for the proof of Theorem B, that come from Proposition 3.

5.3.1. PROPERTIES OF $H^1(\Omega, \sigma)$

Lemma 9 *Under the assumptions of Theorem B, $H^1(\Omega, \sigma)$ has a Poincaré inequality. That is, for all $u \in H^1(\Omega, \sigma)$,*

$$\int_\Omega (u(x) - (u)_\sigma)^2 \sigma(x)dx \leq C\int_\Omega |\nabla u(x)|^2 \sigma(x)dx, \tag{20}$$

where

$$(u)_\sigma = \int_\Omega u(x)\sigma(x)dx.$$

Further, $C^\infty(\overline{\Omega})$ is dense in $H^1(\Omega, \sigma)$ in the $H^1(\Omega, \sigma)$ norm.

Lemma 9 follows by noting that the desired properties hold for the weighted Sobolev space with weight $\rho(x) = d(x, \partial\Omega)$, and using Proposition 3 to show that they are then inherited by $H^1(\Omega, \sigma)$. See Appendix A.3 (in particular, Lemmas 18 and 19) for detailed proofs.

Note that since inequality (14) holds for all $u \in C^\infty(\bar{\Omega})$, density of this space in $H^1(\Omega, \sigma)$ makes $\langle u, g - f \rangle$ well defined for all $u \in H^1(\Omega, \sigma)$.

5.3.2. PROOF SKETCH FOR THEOREM B

We now leverage the properties of $H^1(\Omega, \sigma)$ to prove Theorem B. Lemma 10 below summarizes the required properties of the corresponding functional.

Observing that the functional in $(\widetilde{GP_\lambda})$ is invariant under the map $u \mapsto u + c$, we may restrict without loss of generality to the Hilbert space

$$\bar{H}^1(\Omega, \sigma) = \{u \in H^1(\Omega, \sigma) \mid \int_\Omega u(x)\sigma(x)dx = 0\},$$

where the norm is the same as the one for $H^1(\Omega, \sigma)$. Define $\tilde{J} : \bar{H}^1(\Omega, \sigma) \to \mathbb{R}$ as

$$\tilde{J}(u) = \langle u, g - f \rangle + \frac{\lambda}{2} \int_\Omega (|\nabla u| - 1)^2_+ \sigma(x)dx. \tag{21}$$

It is clear that $\sup(\widetilde{GP_\lambda}) = - \inf_{u \in \bar{H}^1(\Omega, \sigma)} \tilde{J}(u, \nabla u)$.

Lemma 10 *Under the assumptions of Theorem B, $\tilde{J}$ is coercive on $\bar{H}^1(\Omega, \sigma)$ (i.e. $\tilde{J}(u) \to +\infty$ if $\|u\|_{H^1(\Omega, \sigma)} \to +\infty$). Moreover, $\tilde{J}$ is convex, proper, and continuous, and hence also weakly lower semi-continuous.*

This is proved in Appendix A.3, and provides the results needed for the proof of Theorem B.

Proof sketch for Theorem B Let us note that the hypotheses of Theorem A are subsumed by those of Theorem B, so we need only prove statements 2 through 4. The proof of statement 2 of Theorem B follows from Lemma 10 and the direct method of the calculus of variations. The equivalence of the optimal values in statement 3 follows since the functionals in (GP_λ) and $(\widetilde{GP_\lambda})$ agree on $C^\infty(\bar{\Omega})$ and are continuous on both $H^1(\Omega)$ and $H^1(\Omega, \sigma)$, respectively. Finally, statement 4 of Theorem B follows in a similar manner to Proposition 8, except that the existence of a solution to $(\widetilde{GP_\lambda})$ allows us to make the inequality an equality; inverting the map $\xi \mapsto (|\xi| - 1)_+\xi/|\xi|$ for $|\xi| > 1$ then provides the second equality of (11). See Appendix A.3 for a detailed proof. ■

6. Discussion and Conclusion

6.1. Discussion

In this section we provide some intuition on how Theorems A and B might explain the performance of WGAN-GP. We also discuss how our work could be used to create new methods for approximately solving large scale congested transport problems.

We first observe that in congested transport concentration of data is penalized, thus encouraging diverse generated distributions. This may explain how WGAN-GP avoid the strong averaging effect

of the Wasserstein 1 distance which was observed experimentally in [23]. This effect was elucidated in [23] by an inequality which shows that the average Wasserstein 1 distance between minibatches of a given dataset is larger than the average distance between a minibatch and the global mean of the dataset. We expect that under congested transport this inequality would no longer hold, as the transport to the global mean would be expensive since all mass converges to a single point, thus inducing congestion.

Second, our formula (11) may explain why generators trained with WGAN-GP converge faster than those trained with weight clipping, as observed in [12]. In the case of weight clipping, the gradient of the critic is bounded by a fixed constant. On the other hand, formula (11) shows that the gradient of the optimal critic for WGAN-GP depends not only on the direction of optimal flow, but also on the momentum via its dependency on $\mathbf{w}_{Q_0}$ (see (7)). Thus, data points with a long distance to travel under the optimal traffic flow would have a high speed under gradient descent on u_0. We expect that this would lead to faster convergence of generators trained with WGAN-GP.

We emphasize that the intuition given here is based on our results but is not rigorously proven. However, it points to the possibility for future analysis of WGAN-GP based on congested transport theory.

We leave unexamined some important problems in this area, including how our results interact with issues that arise in actual implementations. These include the constrained class of functions over which WGAN-GP are optimized (an issue considered in [3]), issues with approximation of the integrals in (3) by sampling, and the fact that (3) is not solved to completion before updating the generated distribution μ. We also anticipate that there will be practical applications of Theorem **A** to the training of WGANs, and we plan to address these in a follow up paper.

Beyond applications to WGANs, we note that the documentation for [10] states that existing efficient algorithms for computing the Wasserstein 1 distance are not suitable for large datasets, and that critics trained with WGAN-GP are one recommended approach for treating such problems. Theorem **A**, presented here, quantifies the relationship between the congested transport cost computed by WGAN-GP and the Wasserstein 1 distance. As well, our discovery points to the possibility of numerically solving congested transport problems by suitable variants of WGAN-GP.

6.2. Conclusion

We have shown, under weak assumptions on μ and ν, that the value of the optimization problem for WGAN-GP with one-sided penalty is equal to the minimal cost of moving μ to ν under a congested transport model. This cost was shown to be strictly larger than $W_1(\mu, \nu)$ for all $\lambda > 0$ provided $\mu \neq \nu$, with an error that is at least $O(\lambda^{-1})$ as $\lambda \to \infty$. We have also proved, under slightly stronger assumptions on μ and ν, that there are solutions to this problem in an appropriate weighted Sobolev space. The gradients of these solutions are shown to encode more information on the optimal transport of mass than do the gradients of a standard Kantorovich potential (i.e. time averaged momentum, as opposed to just direction of transport). We have also shown that an approximate version of this relationship holds for approximate solutions. This work provides theoretical foundations for understanding the performance of WGAN-GP.

Acknowledgments

References

[1] Cem Anil, James Lucas, and Roger Grosse. Sorting out Lipschitz function approximation. In *International Conference on Machine Learning*, pages 291–301. PMLR, 2019.

[2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223. PMLR, 2017.

[3] Gérard Biau, Maxime Sangnier, and Ugo Tanielian. Some theoretical insights into Wasserstein GANs. *Journal of Machine Learning Research*, 2021.

[4] Lorenzo Brasco and Guillaume Carlier. Congested traffic equilibria and degenerate anisotropic PDEs. *Dynamic Games and Applications*, 3(4):508–522, 2013.

[5] Lorenzo Brasco and Mircea Petrache. A continuous model of transportation revisited. *Journal of mathematical sciences*, 196(2):119–137, 2014.

[6] Lorenzo Brasco, Guillaume Carlier, and Filippo Santambrogio. Congested traffic dynamics, weak flows and very degenerate elliptic equations. *Journal de mathématiques pures et appliquées*, 93(6):652–671, 2010.

[7] Guillaume Carlier, Chloé Jimenez, and Filippo Santambrogio. Optimal transportation with traffic congestion and Wardrop equilibria. *SIAM Journal on Control and Optimization*, 47(3): 1330–1350, 2008.

[8] David E Edmunds and Bohumir Opic. Weighted Poincaré and Friedrichs inequalities. *Journal of the London Mathematical Society*, 2(1):79–96, 1993.

[9] Ivar Ekeland and Roger Temam. *Convex analysis and variational problems*. SIAM, 1999.

[10] Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, et al. Pot: Python optimal transport. *Journal of Machine Learning Research*, 22(78):1–8, 2021.

[11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.

[12] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems*, pages 5767–5777, 2017.

[13] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018.

[14] Alois Kufner. *Weighted Sobolev spaces*, volume 31. John Wiley & Sons Incorporated, 1985.

[15] Alois Kufner and Bohumír Opic. How to define reasonably weighted Sobolev spaces. *Commentationes Mathematicae Universitatis Carolinae*, 25(3):537–554, 1984.

[16] Karol Kurach, Mario Lučić, Xiaohua Zhai, Marcin Michalski, and Sylvain Gelly. A large-scale study on regularization and normalization in GANs. In *International Conference on Machine Learning*, pages 3581–3590. PMLR, 2019.

[17] Sebastian Lunz, Ozan Öktem, and Carola-Bibiane Schönlieb. Adversarial regularizers in inverse problems. In *Advances in Neural Information Processing Systems*, pages 8507–8516, 2018.

[18] Anton Mallasto, Guido Montúfar, and Augusto Gerolin. How well do WGANs estimate the Wasserstein metric? *arXiv preprint arXiv:1910.03875*, 2019.

[19] Tristan Milne, Étienne Bilocq, and Adrian Nachman. Trust the Critics: Generatorless and multipurpose WGANs with initial convergence guarantees. *arXiv preprint arXiv:2111.15099*, 2021.

[20] Henning Petzka, Asja Fischer, and Denis Lukovnikov. On the regularization of Wasserstein GANs. In *International Conference on Learning Representations*, 2018.

[21] Thomas Pinetz, Daniel Soukup, and Thomas Pock. On the estimation of the Wasserstein distance in generative models. In *German Conference on Pattern Recognition*, pages 156–170. Springer, 2019.

[22] Filippo Santambrogio. *Optimal Transport for Applied Mathematicians: Calculus of Variations, PDEs, and Modeling*, volume 87. Birkhäuser, 2015.

[23] Jan Stanczuk, Christian Etmann, Lisa Maria Kreusser, and Carola-Bibiane Schönlieb. Wasserstein GANs work because they fail (to approximate the Wasserstein distance). *arXiv preprint arXiv:2103.01678*, 2021.

[24] John Glen Wardrop. Road paper. Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers*, 1(3):325–362, 1952.

Appendix A. Proofs

In this appendix we will provide detailed proofs of all mathematical results presented in the main body of the paper. The proofs here are organized by the section in which the corresponding result appeared.

A.1. Proofs of results from Section 5.1

Proof of Lemma 1 The measure σ from [12] assigns to a measurable set E the value

$$\sigma(E) = \int_0^1 \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} 1_E((1-t)x + ty)f(x)g(y)dxdydt.$$

Here we have extended the densities f and g by zero to $\mathbb{R}^d$. For y and t fixed with $t \neq 1$, we set $z = (1-t)x + ty$. Applying this change of variable to the integral with respect to x, we obtain

$$\sigma(E) = \int_0^1 \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} 1_E(z) f\left(\frac{z - ty}{1 - t}\right) g(y)(1-t)^{-d} dz\, dy\, dt,$$

$$= \int_E \int_0^1 \int_\Omega f\left(\frac{z - ty}{1 - t}\right) g(y)(1-t)^{-d} dy\, dt\, dz,$$

the second line being obtained by an application of Fubini's theorem and recalling that g is supported in Ω. This proves (12). ∎

Proof of Lemma 2 We write

$$\sigma(z) = \int_0^{1/2} \int_\Omega f\left(\frac{z - ty}{1 - t}\right) g(y)(1-t)^{-d} dy\, dt + \int_{1/2}^1 \int_\Omega f\left(\frac{z - ty}{1 - t}\right) g(y)(1-t)^{-d} dy\, dt.$$

For fixed z and $t \neq 1$, apply the change of variable $y' = \frac{z - ty}{1 - t}$ to the second integral. This gives

$$\int_{1/2}^1 \int_\Omega f\left(\frac{z - ty}{1 - t}\right) g(y)(1-t)^{-d} dy\, dt = \int_{1/2}^1 \int_\Omega f(y') g\left(\frac{z - (1-t)y'}{t}\right) t^{-d} dy'\, dt.$$

As such,

$$\sigma(z) \leq \int_0^{1/2} \int_\Omega \|f\|_{L^\infty(\Omega)} g(y)(1-t)^{-d} dy\, dt + \int_{1/2}^1 \int_\Omega \|g\|_{L^\infty(\Omega)} f(y') t^{-d} dy'\, dt,$$

$$= C(\|f\|_{L^\infty(\Omega)} + \|g\|_{L^\infty(\Omega)}),$$

which proves the claim. ∎

Lemmas 11 and 13 below prove Proposition 3.

Lemma 11 *Suppose that $\inf_{x \in \Omega} f(x) = \epsilon > 0$. Then there exists C, depending only on Ω, such that*

$$\sigma(x) \geq C\epsilon \operatorname{dist}(x, \partial\Omega). \tag{22}$$

Proof Let $x \in \Omega$ with $\operatorname{dist}(x, \partial\Omega) > 0$. Then for $t \in [0, 1)$ and $y \in \Omega$,

$$\left| x - \frac{x - ty}{1 - t} \right| = \frac{t}{1 - t}|x - y| \leq \frac{t}{1 - t} \operatorname{diam}(\Omega),$$

and $\frac{t}{1-t} \operatorname{diam}(\Omega) \leq \operatorname{dist}(x, \partial\Omega)$ if and only if

$$t \leq \frac{\operatorname{dist}(x, \partial\Omega)}{\operatorname{diam}(\Omega) + \operatorname{dist}(x, \partial\Omega)} =: t_0.$$

As such, for all $t \in [0, t_0]$ and $y \in \Omega$, $\frac{x - ty}{1 - t} \in \Omega$, and hence the lower bound on f applies. Thus,

$$
\begin{aligned}
\sigma(x) &= \int_0^1 \int_\Omega f\left(\frac{x - ty}{1 - t}\right) g(y)(1 - t)^{-d} dy \, dt, \\
&\geq \int_0^{t_0} \int_\Omega \epsilon g(y)(1 - t)^{-d} dy \, dt, \\
&= \epsilon \int_0^{t_0} (1 - t)^{-d} dt, \\
&\geq \epsilon t_0,
\end{aligned}
$$

where the last line follows from noting that $(1 - t)^{-d} \geq 1$ for all $t \in [0, 1)$. Hence

$$
\begin{aligned}
\sigma(x) &\geq \epsilon t_0, \\
&= \epsilon \frac{\operatorname{dist}(x, \partial\Omega)}{\operatorname{diam}(\Omega) + \operatorname{dist}(x, \partial\Omega)}, \\
&\geq C\epsilon \operatorname{dist}(x, \partial\Omega),
\end{aligned}
$$

for

$$
C = \frac{1}{\operatorname{diam}(\Omega) + \sup_{x \in \Omega} \operatorname{dist}(x, \partial\Omega)}.
$$

∎

To establish an upper bound on σ in terms of $\operatorname{dist}(x, \partial\Omega)$ we need a basic lemma about convex sets. This result is standard, but we include it here for the convenience of the reader.

Lemma 12 *If Ω is convex, $x \in \Omega$ and $y \in \partial\Omega$ satisfying*

$$
|x - y| = \operatorname{dist}(x, \partial\Omega),
$$

then

$$
\langle y - x, y \rangle = \sup_{z \in \Omega} \langle y - x, z \rangle. \tag{23}
$$

Proof Since $y \in \partial\Omega$ we immediately have

$$
\langle y - x, y \rangle \leq \sup_{z \in \Omega} \langle y - x, z \rangle.
$$

Now suppose that y does not obtain the supremum in (23). Then there exists $z \in \Omega$ such that

$$
\langle y - x, y \rangle < \langle y - x, z \rangle.
$$

Observe that for all $t > 1$, the point $(1 - t)z + ty \notin \Omega^o$ (here Ω^o is the interior of Ω), since otherwise we would have $y \in \Omega^o$ by virtue of Ω being convex.

17

Next we assert that for all $t > 1$ and small enough, the point $(1-t)z + ty$ is strictly closer to x than y is. Indeed,

$$
\begin{aligned}
|(1-t)z + ty - x|^2 &= |(1-t)z - (1-t)y - (x-y)|^2, \\
&= (1-t)^2|z-y|^2 + 2(1-t)\langle z - y, y - x\rangle + |x-y|^2, \\
&= (t-1)((t-1)|z-y|^2 - 2\langle z - y, y - x\rangle) + |x-y|^2.
\end{aligned}
$$

Since $\langle z - y, y - x\rangle > 0$, the assertion is proven. For $t > 1$ and small enough there must therefore exist a point on the segment connecting x and $(1-t)z + ty$ which is in $\partial\Omega$ (since $(1-t)z + ty \notin \Omega^o$) and is closer to x than y, contradicting the assumption that y is the closest point to x in $\partial\Omega$. ∎

We can now establish the desired upper bound provided Ω is convex and g is zero in a neighbourhood of the boundary of Ω.

Lemma 13 *If Ω is convex and open, $f, g \in L^\infty(\Omega)$, and $g = 0$ in a neighbourhood of the boundary of Ω, then there exists a constant C such that for all $x \in \Omega$,*

$$
\sigma(x) \leq C \operatorname{dist}(x, \partial\Omega).
$$

The idea of the proof is to partition Ω into two sets by a certain half-plane, one set in which g is zero, and another set in which the time integral in (12) can be restricted to a sub-interval $[0, t_0]$ where t_0 is related to $\operatorname{dist}(x, \partial\Omega)$.

Proof Let $x \in \Omega$, and let $y \in \partial\Omega$ such that

$$
|x - y| = \operatorname{dist}(x, \partial\Omega).
$$

Set $\theta =: \langle y - x, y\rangle$ and note that by Lemma 12, the condition $\langle y - x, z\rangle > \theta$ implies that $z \in \Omega^c$. For $\delta > 0$ to be determined, set

$$
H^-(x) = \{z \in \Omega \mid \langle y - x, z\rangle \leq \theta - \delta\}, H^+(x) = \{z \in \Omega \mid \langle y - x, z\rangle > \theta - \delta\}.
$$

Since $H^+(x)$ and $H^-(x)$ partition Ω, we can write

$$
\sigma(x) = \int_{H^-(x)} \int_0^1 f\left(\frac{x - tz}{1 - t}\right) g(z)(1-t)^{-d} dt dz + \int_{H^+(x)} \int_0^1 f\left(\frac{x - tz}{1 - t}\right) g(z)(1-t)^{-d} dt dz.
$$

Since $g = 0$ in a neighbourhood of $\partial\Omega$, there exists a constant δ_0 such that if $g(z) > 0$, then $\operatorname{dist}(z, \partial\Omega) > \delta_0$. We assert that the condition $\delta \leq \delta_0 \operatorname{dist}(x, \partial\Omega)$ then implies that $g(z) = 0$ on $H^+(x)$. Indeed, if $z \in H^+(x)$ then

$$
\langle z + \frac{\delta}{\operatorname{dist}(x, \partial\Omega)^2}(y - x), y - x\rangle > \theta - \delta + \delta = \theta.
$$

Hence $z + \frac{\delta}{\operatorname{dist}(x, \partial\Omega)^2}(y - x) \in \Omega^c$, indicating that

$$
\operatorname{dist}(z, \partial\Omega) \leq \frac{\delta}{\operatorname{dist}(x, \partial\Omega)} \leq \delta_0,
$$

and hence $g(z) = 0$, proving the assertion. Therefore, if $\delta \leq \delta_0 \operatorname{dist}(x, \partial\Omega)$, we have

$$\sigma(x) = \int_{H^-(x)} \int_0^1 f\left(\frac{x - tz}{1 - t}\right) g(z)(1 - t)^{-d} dt dz.$$

We next determine, given $z \in H^-(x)$, for what range of $t \in [0, 1]$ we have

$$\frac{x - tz}{1 - t} \in \Omega.$$

Again we use Lemma 12. For $z \in H^-(x)$ we compute

$$\begin{aligned}
\langle \frac{x - tz}{1 - t}, y - x \rangle &= \frac{1}{1 - t}\langle x, y - x \rangle - \frac{t}{1 - t}\langle z, y - x \rangle, \\
&\geq \frac{1}{1 - t}(\theta - \operatorname{dist}(x, \partial\Omega)^2) - \frac{t}{1 - t}(\theta - \delta), \\
&= \theta + \frac{1}{1 - t}(t\delta - \operatorname{dist}(x, \partial\Omega)^2).
\end{aligned}$$

Thus, if $t \in (\frac{\operatorname{dist}(x,\partial\Omega)^2}{\delta}, 1]$ and $z \in H^-(x)$, we have $\frac{x-tz}{1-t} \in \Omega^c$.

Under the assumption that $\operatorname{dist}(x, \partial\Omega) \leq \frac{\delta_0}{2}$, we now select

$$\delta := \delta_0 \operatorname{dist}(x, \partial\Omega).$$

With this choice, the interval $(\frac{\operatorname{dist}(x,\partial\Omega)^2}{\delta}, 1]$ is non-empty since $\frac{\operatorname{dist}(x,\partial\Omega)^2}{\delta} = \frac{\operatorname{dist}(x,\partial\Omega)}{\delta_0} \leq \frac{1}{2}$. Since $\delta \leq \delta_0 \operatorname{dist}(x, \partial\Omega)$, we use our work above to conclude that

$$\begin{aligned}
\sigma(x) &= \int_{H^-(x)} \int_0^{\frac{\operatorname{dist}(x,\partial\Omega)}{\delta_0}} f\left(\frac{x - tz}{1 - t}\right) g(z)(1 - t)^{-d} dt dz, \\
&\leq \|f\|_{L^\infty(\Omega)} \int_{H^-(x)} \int_0^{\frac{\operatorname{dist}(x,\partial\Omega)}{\delta_0}} (1 - t)^{-d} g(z) dt dz.
\end{aligned}$$

Noting that $(1 - t)^{-d}$ is an increasing function on $[0, 1)$ and $\frac{\operatorname{dist}(x,\partial\Omega)}{\delta_0} \leq \frac{1}{2}$, we get

$$\sigma(x) \leq \|f\|_{L^\infty(\Omega)} \frac{2^d}{\delta_0} \operatorname{dist}(x, \partial\Omega).$$

For x with $\operatorname{dist}(x, \partial\Omega) > \frac{\delta_0}{2}$ we use Lemma 2 to conclude

$$\begin{aligned}
\sigma(x) &\leq \|\sigma\|_{L^\infty(\Omega)}, \\
&\leq \frac{2\|\sigma\|_{L^\infty(\Omega)}}{\delta_0} \operatorname{dist}(x, \partial\Omega).
\end{aligned}$$

Hence, for all $x \in \Omega$,

$$\sigma(x) \leq \max\left(\frac{2}{\delta_0}\|\sigma\|_{L^\infty(\Omega)}, \frac{2^d}{\delta_0}\|f\|_{L^\infty(\Omega)}\right) \operatorname{dist}(x, \partial\Omega).$$

$\blacksquare$

A.2. Proofs of results from Section 5.2

Proof of Lemma 4 For $u \in C^\infty(\bar{\Omega})$, we have

$$|\langle u, g - f \rangle| = |\int_\Omega u(x)(g(x) - f(x))dx|,$$

$$= |\int_\Omega \int_\Omega \int_0^1 \nabla u((1-t)x + ty) \cdot (y - x)dt f(x)g(y)dxdy|,$$

$$\leq \operatorname{diam}(\Omega) \int_\Omega \int_\Omega \int_0^1 |\nabla u((1-t)x + ty)| f(x)g(y)dtdxdy,$$

$$= \operatorname{diam}(\Omega) \int_\Omega |\nabla u(x)|\sigma(x)dx.$$

Thus, via Cauchy-Schwarz,

$$|\langle u, g - f \rangle| \leq \operatorname{diam}(\Omega) \left(\int_\Omega |\nabla u|^2 \sigma(x)dx \right)^{1/2}.$$

By density of $C^\infty(\bar{\Omega})$ in $H^1(\Omega)$, (14) therefore holds for all $u \in H^1(\Omega)$. As such, for $u \in H^1(\Omega)$,

$$J(u, \nabla u) \geq - \operatorname{diam}(\Omega) \left(\int_\Omega |\nabla u|^2 \sigma(x)dx \right)^{1/2} + \frac{\lambda}{2} \int_\Omega (|\nabla u| - 1)_+^2 \sigma(x)dx,$$

$$\geq - \operatorname{diam}(\Omega) \left(\int_\Omega (|\nabla u| - 1)_+^2 \sigma(x)dx \right)^{1/2} + \frac{\lambda}{2} \int_\Omega (|\nabla u| - 1)_+^2 \sigma(x)dx - C,$$

$$= \frac{\lambda}{2} \left(\left(\int_\Omega (|\nabla u| - 1)_+^2 \sigma(x)dx \right)^{1/2} - \frac{\operatorname{diam}(\Omega)}{\lambda} \right)^2 - C,$$

$$> -C,$$

where the second inequality is obtained by applying the triangle inequality twice. This proves that $\inf_{u \in H^1(\Omega)} J(u, \nabla u)$ (and hence $\sup (GP_\lambda)$) is finite. ∎

To established the duality result in Proposition 5, we need some simple properties of the functional J, which we prove here.

Lemma 14 *The functional J defined in (13) is convex and continuous.*

Proof Define $F : H^1(\Omega) \to \mathbb{R}$ and $G : L^2(\Omega; \mathbb{R}^d) \to \mathbb{R}$ as

$$F(u) = \langle u, g - f \rangle, \quad G(\mathbf{p}) = \frac{\lambda}{2} \int_\Omega (|\mathbf{p}| - 1)_+^2 \sigma(x)dx,$$

so that $J(u, \mathbf{p}) = F(u) + G(\mathbf{p})$. To see that J is convex and continuous, we claim that F and G are both convex and continuous on their respective domains. For F these claims are immediate since F is a continuous linear operator. For G they require proof. Note that G can be written

$$G(\mathbf{p}) = \int_\Omega h(x, \mathbf{p}(x))dx,$$

where $h : \Omega \times \mathbb{R}^d \to \mathbb{R}$ is given by

$$h(x, \xi) = \frac{\lambda}{2}(|\xi| - 1)_+^2 \sigma(x). \tag{24}$$

Note that for fixed x, h is convex in ξ because it is the composition of a convex function $(\xi \mapsto |\xi|)$ with a non-decreasing convex function $(z \mapsto \frac{\lambda}{2}(z-1)_+^2 \sigma(x))$. Convexity of h then implies convexity of G.

To show continuity of G, observe that for $\mathbf{p}_1, \mathbf{p}_2 \in L^2(\Omega; \mathbb{R}^d)$,

$$G(\mathbf{p}_1) - G(\mathbf{p}_2) = \frac{\lambda}{2}\int_\Omega (|\mathbf{p}_1| - 1)_+ + (|\mathbf{p}_2| - 1)_+)(|\mathbf{p}_1| - 1)_+ - (|\mathbf{p}_2| - 1)_+)\sigma(x)dx,$$

$$\leq \frac{\lambda}{2}\|\sigma\|_{L^\infty(\Omega)}(\|\mathbf{p}_1\|_{L^2(\Omega;\mathbb{R}^d)} + \|\mathbf{p}_2\|_{L^2(\Omega;\mathbb{R}^d)})$$

$$\times \left(\int_\Omega ((|\mathbf{p}_1| - 1)_+ - (|\mathbf{p}_2| - 1)_+)^2 \, dx\right)^{1/2},$$

$$\leq \frac{\lambda}{2}\|\sigma\|_{L^\infty(\Omega)}(\|\mathbf{p}_1\|_{L^2(\Omega;\mathbb{R}^d)} + \|\mathbf{p}_2\|_{L^2(\Omega;\mathbb{R}^d)})\|\mathbf{p}_1 - \mathbf{p}_2\|_{L^2(\Omega;\mathbb{R}^d)},$$

where in the last line we have used the fact that $\xi \mapsto (|\xi| - 1)_+$ is 1-Lipschitz. Swapping the roles of $\mathbf{p}_1$ and $\mathbf{p}_2$, we obtain that G is continuous. ∎

The same proof, with minor modifications, establishes the following related result.

Lemma 15 *The map $\tilde{G} : L^2(\Omega, \sigma; \mathbb{R}^d) \to \mathbb{R}$ given by*

$$\tilde{G}(\mathbf{p}) = \frac{\lambda}{2}\int_\Omega (|\mathbf{p}| - 1)_+^2 \sigma(x)dx \tag{25}$$

is continuous and convex.

Proof of Proposition 5 We apply Theorem 4.1 of Chapter 3 of [9], which states that if

i. J is convex,

ii. $\inf_{u \in H^1(\Omega)} J(u, \nabla u)$ is finite, and

iii. there exists $u_0 \in H^1(\Omega)$ such that $J(u_0, \nabla u_0) < \infty$ with the function $\mathbf{p} \mapsto J(u_0, \mathbf{p})$ being continuous at ∇u_0,

then with J^* as the Legendre dual of J,

$$\sup_{\mathbf{w} \in L^2(\Omega;\mathbb{R}^d)} -J^*(\nabla^*\mathbf{w}, -\mathbf{w}) = \inf_{u \in H^1(\Omega)} J(u, \nabla u),$$

and the problem on the left hand side has at least one solution. By Lemma 16, this is equivalent to (15) and existence of a solution to (BP_λ). Note that this solution must be unique by strict convexity of the functional in (BP_λ). Thus if we can verify points i - iii we are done, and these are shown in Lemmas 4 and 14. ∎

Lemma 16 *Identifying $L^2(\Omega; \mathbb{R}^d)$ as its own dual space, we have that*

$$J^*(\nabla^*\mathbf{w}, -\mathbf{w}) = \begin{cases} \int_\Omega H(x, |\mathbf{w}|(x))dx & \text{if} \quad \nabla \cdot \mathbf{w} = f - g, \\ +\infty & \text{else}. \end{cases}$$

Proof of Lemma 16 We have

$$J^*(u^*, \mathbf{p}^*) = F^*(u^*) + G^*(\mathbf{p}^*).$$

Since F is linear, the definition of the Legendre dual gives that

$$F^*(u^*) = \begin{cases} 0 & u^* = \langle \cdot, g - f \rangle, \\ +\infty & \text{else}, \end{cases} = 1_{g-f}(u^*).$$

To calculate $G^*(\mathbf{p}^*)$, we use Proposition 2.1 from Chapter 9 of [9], which relies on the measurable selection theorem. In this context we write

$$G(\mathbf{p}) = \int_\Omega h(x, \mathbf{p}(x))dx,$$

where h is as in (24). We verify the hypotheses of this proposition, which are

  i. that Ω is a bounded open subset of $\mathbb{R}^d$

  ii. that h is a non-negative normal integrand (see Definition 1.1, Chapter 8, [9]), and

  iii. that there exists $\mathbf{p} \in L^\infty(\Omega; \mathbb{R}^d)$ with $G(\mathbf{p}) < +\infty$.

The third point is clear for $\mathbf{p} = 0$. The second point follows because h is a Carathéodory function, which are proven to be normal integrands in Proposition 1.1, Chapter 8, [9]; recall that a function $h : \Omega \times \mathbb{R}^d \to \mathbb{R}$ is said to be a Carathéodory function if

- for almost all $x \in \Omega$, $h(x, \cdot)$ is continuous on $\mathbb{R}^d$, and

- for almost all $\xi \in \mathbb{R}^d$, $h(\cdot, \xi)$ is measurable on Ω.

Both of these clearly hold for h given in (24). The conclusion of Proposition 2.1 from Chapter 9 of [9] is that

$$G^*(\mathbf{p}^*) = \int_\Omega h^*(x, \mathbf{p}^*(x))dx,$$

for h^* the partial dual

$$h^*(x, \xi^*) = \sup_{\xi \in \mathbb{R}^d} \xi^* \cdot \xi - h(x, \xi).$$

An elementary calculation gives that

$$h^*(x, \xi^*) = H(x, |\xi^*|),$$

where H is as in (9). Relabelling the variable as $\mathbf{w}$ for consistency, we have

$$J^*(\nabla^* \mathbf{w}, -\mathbf{w}) = 1_{g-f}(\nabla^* \mathbf{w}) + \int_\Omega H(x, |\mathbf{w}|(x))dx.$$

An infinite value is obtained for the first term unless $\mathbf{w}$ is such that

$$\int_\Omega \mathbf{w} \cdot \nabla u \, dx = \int_\Omega u(g-f)dx$$

for all $u \in H^1(\Omega)$. Writing this requirement as $\nabla \cdot \mathbf{w} = f - g$, the proof is complete. ∎

Proof of Proposition 6 Here we will follow the strategy outlined in the sketch of the proof in the main text. Let Q be admissible for (CP_λ). It is known (see [22], Section 4.2.3) that

$$|\mathbf{w}_Q| \leq i_Q, \tag{26}$$

where $|\mathbf{w}_Q|$ is the total variation measure of $\mathbf{w}_Q$. Then by (26) and the definition of (CP_λ) we know that $|\mathbf{w}_Q| \ll \mathcal{L}_d$ and $\mathbf{w}_Q \in L^2(\Omega; \mathbb{R}^d)$. Take $u \in C^\infty(\overline{\Omega})$, and observe that

$$\begin{aligned}
\int_\Omega \nabla u \cdot \mathbf{w}_Q dx &= \int_\mathcal{C} \int_0^1 \nabla u(\omega(t)) \cdot \omega'(t)dt dQ, \\
&= \int_\mathcal{C} (u(\omega(1)) - u(\omega(0)))dQ, \\
&= \int_\Omega u(g-f)dx,
\end{aligned}$$

where in the last line we have used the definition of $\mathcal{Q}(\mu, \nu)$. Thus, by density of $C^\infty(\overline{\Omega})$ in $H^1(\Omega)$, we have that for all $u \in H^1(\Omega)$,

$$\int_\Omega \nabla u \cdot \mathbf{w}_Q dx = \langle u, g - f \rangle.$$

Hence $\mathbf{w}_Q$ is admissible in (BP_λ). Moreover, (26) together with the monotonicity of the integrand in (CP_λ) gives

$$\int_\Omega H(x, i_Q(x))dx \geq \int_\Omega H(x, |\mathbf{w}_Q|(x))dx.$$

This establishes $\inf(CP_\lambda) \geq \inf(BP_\lambda)$. Now let $\mathbf{w}_0$ be a solution to (BP_λ). Since $C^\infty(\overline{\Omega}) \subset H^1(\Omega)$, we have

$$\int_\Omega \nabla u \cdot \mathbf{w}_0 dx = \int_\Omega u(g-f)dx$$

for all $u \in C^\infty(\overline{\Omega})$, which is the hypothesis for Theorem 4.10 of [22]. As such, there exists $Q_0 \in \mathcal{P}(\mathcal{C})$ such that $(e_0)_\# Q_0 = \mu, (e_1)_\# Q_0 = \nu$, and

$$i_{Q_0} = |\mathbf{w}_{Q_0}| \leq |\mathbf{w}_0|. \tag{27}$$

Hence $i_{Q_0} \ll \mathcal{L}_d$ and $i_{Q_0} \in L^2(\Omega)$, so Q_0 is admissible for (CP_λ). Further, (27) implies that

$$\int_\Omega H(x, i_{Q_0}(x))dx \leq \int_\Omega H(x, |\mathbf{w}_0|(x))dx,$$

and thus $\inf(CP_\lambda) \le \inf(BP_\lambda)$, establishing (16) and that Q_0 is a solution to (CP_λ).

To prove the final claim, let Q_0 be optimal in (CP_λ). Then the inequality (26), together with (16), shows that $\mathbf{w}_{Q_0}$ is optimal for (BP_λ). Recalling from Proposition 5 that (BP_λ) has a unique minimizer, we get the first equality of (17). To get the second equality we note that it is implied by optimality of $\mathbf{w}_{Q_0}$ together with (16) and (26). ∎

Proof of Lemma 7 By Proposition 5 and since $\sigma \in L^\infty(\Omega)$, it is clear that

$$\sup(GP_\lambda) \ge \inf\left\{\int_\Omega \frac{1}{2\lambda\|\sigma\|_{L^\infty(\Omega)}}|\mathbf{w}|^2(x)dx \mid \mathbf{w} \in L^2(\Omega;\mathbb{R}^d), \nabla \cdot \mathbf{w} = \mu - \nu\right\}$$

$$+ \inf\left\{\int_\Omega |\mathbf{w}|(x)dx \mid \mathbf{w} \in L^2(\Omega;\mathbb{R}^d), \nabla \cdot \mathbf{w} = \mu - \nu\right\}. \tag{28}$$

Let us analyse the second problem. Following Chapter 4 of [22], write $\mathcal{M}_{\mathrm{div}}^d(\Omega)$ as the set of vector measures with divergence which is a scalar measure. It is clear that

$$\inf\left\{\int_\Omega |\mathbf{w}|(x)dx \mid \mathbf{w} \in L^2(\Omega;\mathbb{R}^d), \nabla\cdot\mathbf{w} = \mu-\nu\right\} \ge \inf\{|\mathbf{w}|(\Omega) \mid \mathbf{w} \in \mathcal{M}_{\mathrm{div}}^d(\Omega), \nabla\cdot\mathbf{w} = \mu-\nu\},$$

where $|\mathbf{w}|$ is the total variation measure of $\mathbf{w}$, and it is well known (see Chapter 4 of [22] again) that the right hand side is $W_1(\mu, \nu)$. We use Cauchy-Schwarz in (28) to get

$$\sup(GP_\lambda) \ge \frac{1}{2\lambda\|\sigma\|_{L^\infty(\Omega)}\mathrm{Vol}(\Omega)}W_1(\mu,\nu)^2 + W_1(\mu,\nu),$$

whence the conclusion follows immediately. ∎

Lemma 17 provides a lower bound on a function φ that is used to prove Proposition 8.

Lemma 17 *For $\xi^* \in \mathbb{R}^d$, define $\varphi : \Omega \times \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ by*

$$\varphi(x,\xi) = \xi^* \cdot \xi + \frac{\lambda\sigma(x)}{2}(|\xi| - 1)_+^2 + H(x,|\xi^*|).$$

For all $(x,\xi) \in \Omega \times \mathbb{R}^d$ we have

$$\varphi(x,\xi) \ge \frac{1}{2\lambda\|\sigma\|_{L^\infty(\Omega)}}|\xi^* + \lambda\sigma(x)(|\xi| - 1)_+\frac{\xi}{|\xi|}|^2. \tag{29}$$

Proof Note that if $\sigma(x) = 0$ and $\xi^* \ne 0$, $\varphi(x,\xi) = +\infty$, and so (29) holds. Further, if $\sigma(x) = 0$ and $\xi^* = 0$, $\varphi(x,\xi) = 0$, and (29) holds. We therefore proceed assuming that $\sigma(x) \ne 0$. Suppose $|\xi| \ge 1$. Through elementary manipulations one can show that

$$\varphi(x,\xi) = \frac{\lambda\sigma(x)}{2}|\xi - \frac{\xi}{|\xi|} + \frac{\xi^*}{\lambda\sigma(x)}|^2 + \xi^* \cdot \frac{\xi}{|\xi|} + |\xi^*|.$$

Applying the Cauchy-Schwarz inequality,

$$\varphi(x,\xi) \ge \frac{1}{2\lambda\sigma(x)}|\xi^* + \lambda\sigma(x)(|\xi| - 1)_+\frac{\xi}{|\xi|}|^2.$$

For $|\xi| \leq 1$,

$$\varphi(x, \xi) = \xi^* \cdot \xi + \frac{1}{2\lambda\sigma(x)}|\xi^*|^2 + |\xi^*|,$$

$$\geq \frac{1}{2\lambda\sigma(x)}|\xi^*|^2,$$

$$= \frac{1}{2\lambda\sigma(x)}|\xi^* + \lambda\sigma(x)(|\xi| - 1)_+ \frac{\xi}{|\xi|}|^2.$$

Noting that

$$\frac{1}{\sigma(x)} \geq \frac{1}{\|\sigma\|_{L^\infty(\Omega)}},$$

we obtain the inequality (29). ∎

Proof of Proposition 8 For $\xi^* = \mathbf{w}_0(x)$, $\xi = \nabla u(x)$, the inequality (29) gives

$$\nabla u(x) \cdot \mathbf{w}_0(x) + \frac{\lambda}{2}(|\nabla u|(x) - 1)_+^2 \sigma(x) + H(x, |\mathbf{w}_0|(x))$$

$$\geq \frac{1}{2\lambda\|\sigma\|_{L^\infty(\Omega)}}|\mathbf{w}_0(x) + \lambda\sigma(x)(|\nabla u(x)| - 1)_+ \frac{\nabla u(x)}{|\nabla u(x)|}|^2.$$

Integrating this over Ω and using the equality (15) together with $\nabla \cdot \mathbf{w}_0 = f - g$, we get

$$J(u, \nabla u) + \sup(GP_\lambda) \geq \int_\Omega \frac{1}{2\lambda\|\sigma\|_{L^\infty(\Omega)}}|\mathbf{w}_0(x) + \lambda\sigma(x)(|\nabla u(x)| - 1)_+ \frac{\nabla u(x)}{|\nabla u(x)|}|^2 dx,$$

which is (19). ∎

A.3. Proofs of results from Section 5.3

We start by showing that σ being comparable to $\operatorname{dist}(x, \partial\Omega)$ implies that the weighted Sobolev space $H^1(\Omega, \sigma)$ has a Poincaré inequality (Lemma 18) and that $C^\infty(\bar{\Omega})$ is dense in $H^1(\Omega, \sigma)$ (Lemma 19). Together these lemmas prove Lemma 9.

Lemma 18 *Under the hypotheses of Theorem **B** (20) holds for all $u \in H^1(\Omega, \sigma)$.*

Proof Take $\rho(x) = \operatorname{dist}(x, \partial\Omega)$. Then Remark 5.3 of [8] gives that the space $H^1(\Omega, \rho)$ has a Poincaré inequality. Note that Proposition 3 guarantees that $H^1(\Omega, \sigma) = H^1(\Omega, \rho)$ as sets, and hence we have

$$\int_\Omega (u(x) - (u)_\rho)^2 \rho(x)dx \leq C \int_\Omega |\nabla u(x)|^2 \rho(x)dx \tag{30}$$

for all $u \in H^1(\Omega, \sigma)$, where $(u)_\rho = \frac{\int_\Omega u(x)\rho(x)dx}{\int_\Omega \rho(x)dx}$. Now observe that

$$\|u - (u)_\sigma\|_{L^2(\Omega, \sigma)} = \min_{c \in \mathbb{R}}\|u - c\|_{L^2(\Omega, \sigma)}.$$

Indeed, the right hand side is a 1-D optimization problem with optimality condition

$$c = \frac{\int_\Omega u(x)\sigma(x)dx}{\int_\Omega \sigma(x)},$$

the denominator being finite since σ is a probability distribution. As such,

$$
\begin{aligned}
\|u - (u)_\sigma\|_{L^2(\Omega,\sigma)}^2 &\leq \|u - (u)_\rho\|_{L^2(\Omega,\sigma)}^2, \\
&\leq C\|u - (u)_\rho\|_{L^2(\Omega,\rho)}^2, \\
&\leq C\|\nabla u\|_{L^2(\Omega,\rho;\mathbb{R}^d)}^2, \\
&\leq C\|\nabla u\|_{L^2(\Omega,\sigma;\mathbb{R}^d)}^2.
\end{aligned}
$$

In the second line we used Proposition 3, in the third we used the Poincaré inequality for $H^1(\Omega,\rho)$, and in the fourth we used Proposition 3 again. ∎

Lemma 19 *Under the hypotheses of Theorem **B** the space $C^\infty(\bar{\Omega})$ is dense in $H^1(\Omega,\sigma)$.*

Proof Density of $C^\infty(\bar{\Omega})$ in $H^1(\Omega,\rho)$, for $\rho(x) := \mathrm{dist}(x,\partial\Omega)$ and Ω having a Lipschitz boundary is shown, for example, in Theorem 7.2 of [14]. By Proposition 3, σ is comparable to $\rho(x)$, and hence density of $C^\infty(\bar{\Omega})$ in $H^1(\Omega,\rho)$ carries over to $H^1(\Omega,\sigma)$. ∎

Proof of Lemma 10 Let $\tilde{G}$ be as given in (25). For $u \in \bar{H}^1(\Omega,\sigma)$,

$$
\begin{aligned}
\frac{2}{\lambda}\tilde{G}(\nabla u) &= \int_\Omega (|\nabla u|^2 - 2|\nabla u| + 1)\sigma(x)dx - \int_{\{|\nabla u|<1\}} (|\nabla u|^2 - 2|\nabla u| + 1)\sigma(x)dx, \\
&\geq \int_\Omega |\nabla u|^2\sigma(x)dx - 2\left(\int_\Omega |\nabla u|^2\sigma(x)dx\right)^{1/2} - 1.
\end{aligned}
$$

Since $H^1(\Omega,\sigma)$ has a Poincaré inequality, we obtain that for all $u \in \bar{H}^1(\Omega,\sigma)$,

$$
\tilde{G}(\nabla u) \geq C\|u\|_{H^1(\Omega,\sigma)}^2 - \lambda\|u\|_{H^1(\Omega,\sigma)} - \frac{\lambda}{2}.
$$

Since $u \mapsto \langle u, g - f\rangle$ is a continuous linear functional on $H^1(\Omega,\sigma)$, we therefore obtain

$$
\tilde{J}(u) \geq C\|u\|_{H^1(\Omega,\sigma)}^2 - (\lambda + C)\|u\|_{H^1(\Omega,\sigma)} - \frac{\lambda}{2}.
$$

Thus, $\tilde{J}$ is coercive on $\bar{H}^1(\Omega,\sigma)$.

$\tilde{J}$ is obviously proper. It is convex and continuous by Lemma 15. Weak lower semi-continuity then follows since $\tilde{J}$ is convex and continuous. ∎

For clarity of presentation we have separated the portion of Theorem **B** not implied by Theorem **A** into Lemmas 20, 21, and 22; these correspond to statements 2, 3, and 4 of Theorem **B**, respectively.

Lemma 20 *Under the assumptions of Theorem **B**, $(\widetilde{GP_\lambda})$ has a solution.*

Proof Since $\tilde{J}$ (defined in (21)) is proper and coercive (by Lemma 10), the infimum in question is finite. Let $(u_n)_{n=1}^{\infty} \subset \bar{H}^1(\Omega, \sigma)$ be a minimizing sequence, i.e.

$$\lim_{n \to \infty} \tilde{J}(u_n) = \inf_{u \in \bar{H}^1(\Omega, \sigma)} \tilde{J}(u_n).$$

Since this infimum is finite, coercivity of $\tilde{J}$ implies that $(u_n)_{n=1}^{\infty}$ is bounded in $\bar{H}^1(\Omega, \sigma)$. By Banach-Alaoglu there exists a weakly convergent subsequence converging to some $u_0 \in \bar{H}^1(\Omega, \sigma)$, and by weak lower semi-continuity of $\tilde{J}$, u_0 must be a minimizer of $\tilde{J}(u, \nabla u)$ over $H^1(\Omega, \sigma)$. Thus, u_0 is a solution to $(\widetilde{GP_\lambda})$. ■

Lemma 21 *Under the assumptions of Theorem **B**,*

$$\sup (GP_\lambda) = \max (\widetilde{GP_\lambda}). \tag{31}$$

Proof Note that the functionals in (GP_λ) and $(\widetilde{GP_\lambda})$ agree on $C^\infty(\bar{\Omega})$ and are continuous on both $H^1(\Omega)$ and $H^1(\Omega, \sigma)$, respectively. Since $C^\infty(\bar{\Omega})$ is dense in both spaces, (31) then follows. ■

Lemma 22 *Under the assumptions of Theorem **B**, if u_0 solves $(\widetilde{GP_\lambda})$ and Q_0 solves (CP_λ), then* (11) *holds.*

Proof Evaluating inequality (29) at $\xi^* = \mathbf{w}_{Q_0}(x)$, and $\xi = \nabla u_0(x)$ and integrating, we get

$$0 \geq \int_\Omega \frac{1}{2\lambda \|\sigma\|_{L^\infty(\Omega)}} |\mathbf{w}_{Q_0} + \lambda \sigma(x)(|\nabla u_0(x)| - 1)_+ \frac{\nabla u_0(x)}{|\nabla u_0(x)|}|^2 dx,$$

whence the first equality of (11) follows. To obtain the second, note that for almost all x with $\mathbf{w}_{Q_0}(x) \neq 0$, the first equality implies that

$$\nabla u_0(x) = -\alpha(x) \mathbf{w}_{Q_0}(x)$$

for some $\alpha(x) > 0$. Taking the magnitude of both sides of the first equality gives

$$\alpha(x) = \frac{1}{\lambda \sigma(x)} + \frac{1}{|\mathbf{w}_{Q_0}(x)|},$$

as desired.

■