

MIT Open Access Articles

Estimation of Wasserstein distances in the Spiked Transport Model

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Niles-Weed, Jonathan and Rigollet, Philippe. 2022. "Estimation of Wasserstein distances in the Spiked Transport Model." *Bernoulli*, 28 (4).

As Published: 10.3150/21-bej1433

Publisher: Bernoulli Society for Mathematical Statistics and Probability

Persistent URL: <https://hdl.handle.net/1721.1/145840>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Estimation of Wasserstein distances in the Spiked Transport Model

JONATHAN NILES-WEED* AND PHILIPPE RIGOLLET†

New York University and Massachusetts Institute of Technology

Abstract. We propose a new statistical model, the *spiked transport model*, which formalizes the assumption that two probability distributions differ only on a low-dimensional subspace. We study the minimax rate of estimation for the Wasserstein distance under this model and show that this low-dimensional structure can be exploited to avoid the curse of dimensionality. As a byproduct of our minimax analysis, we establish a lower bound showing that, in the absence of such structure, the plug-in estimator is nearly rate-optimal for estimating the Wasserstein distance in high dimension. We also give evidence for a statistical-computational gap and conjecture that any computationally efficient estimator is bound to suffer from the curse of dimensionality.

AMS 2000 subject classifications: Primary 62F99; secondary 62H99.

Key words and phrases: Wasserstein distance; optimal transport; high-dimensional statistics.

1. INTRODUCTION

Optimal transport is an increasingly useful toolbox in various data-driven disciplines such as machine learning [2, 3, 4, 17, 20, 32, 36, 41, 44, 64, 67, 73, 75, 80], computer graphics [31, 54, 78, 79], statistics [1, 7, 18, 19, 21, 43, 49, 51, 55, 65, 70, 71, 72, 76, 82, 91, 94] and the sciences [22, 59, 74, 93]. A core primitive of this toolbox is the computation of Wasserstein distances between probability measures, and a natural statistical question is the estimation of Wasserstein distances from data.

A key object in this endeavor is the empirical measure μ_n associated to μ . It is the empirical measure defined by

$$\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}, \quad X_i \sim \mu \quad \text{i.i.d.}$$

Owing to their flexibility, Wasserstein distances are notoriously hard to estimate in high dimension since in such cases, the empirical distribution is a poor proxy

*Part of this research was conducted while at the Institute for Advanced Study. JNW gratefully acknowledges its support.

†Supported by NSF awards IIS-BIGDATA-1838071, DMS-1712596 and CCF-TRIPODS-1740751; ONR grant N00014-17-1-2147.

for the underlying distribution of interest. Indeed, for d -dimensional distributions, Wasserstein distances between empirical measures generally converge at the slow rate $n^{-1/d}$ [11, 23, 29, 34, 90] and thus suffer from the curse of dimensionality. For example, the following behavior is typical.

PROPOSITION 1. *Let μ be a probability measure on $[-1, 1]^d$. If μ_n is an empirical measure comprising n i.i.d. samples from μ , then for any $p \in [1, \infty)$,*

$$\mathbb{E}W_p(\mu_n, \mu) \leq r_{p,d}(n) := c_p \sqrt{d} \begin{cases} n^{-1/2p} & \text{if } d < 2p \\ n^{-1/2p}(\log n)^{1/p} & \text{if } d = 2p \\ n^{-1/d} & \text{if } d > 2p \end{cases}$$

In many settings, this bound is known to be tight up to logarithmic factors. In fact, the rate in Proposition 1 has been shown to be essentially minimax optimal for the problem of estimating μ in Wasserstein distance [77], using any estimator, not necessarily the empirical distribution μ_n .

Since W_p satisfies the triangle inequality, Proposition 1 readily yields that $W_p(\mu^{(1)}, \mu^{(2)})$ for two probability measures $\mu^{(1)}, \mu^{(2)}$ on $[-1, 1]^d$ can be estimated at the rate $n^{-1/d}$ by the plug-in estimator $W_p(\mu_n^{(1)}, \mu_n^{(2)})$ when $d > 2p$, and a lower bound of the same order can be shown for the plug-in estimator when $\mu^{(1)}$ and $\mu^{(2)}$ are, for example, the uniform measure on $[-1, 1]^d$. However, while the above results give a strong indication that the Wasserstein distance $W_p(\mu^{(1)}, \mu^{(2)})$ itself is also hard to estimate in high dimension, they do not preclude the existence of estimators that are better than $W_p(\mu_n^{(1)}, \mu_n^{(2)})$. Indeed, until recently, the best known lower bound for the problem of estimating the distance itself was of order $n^{-3/2d}$ [28]. A concurrent and independent result [58] closes this gap for $p = 1$ and indicates that estimating the W_1 itself is essentially as hard as estimating the measure itself. Our results show that, in fact, estimating the distance $W_p(\mu^{(1)}, \mu^{(2)})$ is essentially as hard as estimating a measure μ in W_p -distance, for any $p \geq 1$. As a result, any estimator of the distance itself must suffer from the curse of dimensionality.

One goal of *statistical optimal transport* is to develop new models and methods that overcome this curse of dimensionality by leveraging plausible structure in the problem. Early contributions in this direction include assuming smoothness [43, 58, 91] or sparsity [33]. In this work, we propose a new model, called the *spiked transport model*, to formalize the assumption that two distributions differ only on a low dimensional subspace of \mathbb{R}^d . Such an assumption forms the basis of several popular alternatives to Wasserstein distances such as the Sliced Wasserstein distance [69] or random one-dimensional projections [68]. More recently, several numerical methods that exploit a form of low-dimensional structure were proposed together with illustration of their good numerical performance [24, 50, 66].

To exploit the low-dimensional structure of the spiked transport model, we consider a standard method in statistics often called “projection pursuit” [35, 52, 53]. This general method aims to alleviate the challenges of high-dimensional data analysis by considering low-dimensional projections of the dataset that reveal “interesting” features of the data. We show that a suitable instantiation of this method to the present problem, which we call *Wasserstein Projection Pursuit* (WPP), leads to near optimal rates of estimation of the Wasserstein distance in

the spiked transport model and permits to alleviate the curse of dimensionality from which the plug-in estimator suffers.

While our results establish a clear statistical picture, it is unclear how to implement WPP efficiently. An efficient relaxation of this estimator was recently proposed by Paty and Cuturi [66], and a natural question is to analyze its performance in the spiked transport model. Instead of pursuing this direction we bring strong evidence that, in fact, *no* computationally efficient estimator is likely to be able to take advantage of the low-dimensional structure inherent to the spiked transport model. Our computational lower bounds come from the well-established statistical query framework [48]. In particular, they indicate a fundamental trade-off between statistical and computational efficiency [5, 6, 13, 16, 60]: computationally efficient methods to estimate Wasserstein distances, are bound to suffer the curse of dimensionality.

The rest of this paper is organized as follows. We introduce our model and main results in Sections 2 and 3. In Section 4, we define a low-dimensional version of the Wasserstein distance and establish its connection to the spiked transport model. Section 5 proves the equivalence between transportation inequalities and subgaussian concentration properties of the Wasserstein distance. In Section 6 we propose and analyze an estimator for the Wasserstein distance under the spiked transport model. We establish a minimax lower bound in Section 7. Finally, we prove a statistical query bound on the performance of efficient estimators in Section 8. Supplementary proofs and lemmas appear in the appendices.

NOTATION.

We denote by $\|\cdot\|$ the Euclidean norm on \mathbb{R}^d . The symbols $\|\cdot\|_{\text{op}}$ and $\|\cdot\|_{\text{F}}$ denote the operator norm and Frobenius norm, respectively. If X is a random variable on \mathbb{R} , we let $\|X\|_p := (\mathbb{E}|X|^p)^{1/p}$. Throughout, we use c and C to denote positive constants whose value may change from line to line, and we use subscripts to indicate when these constants depend on other parameters. We write $a \lesssim b$ if $a \leq Cb$ holds for a universal positive constant C .

2. MODEL AND METHODS

In this section, we describe the spiked transport model and Wasserstein projection pursuit.

2.1 Wasserstein distances

Given two probability measures μ and ν on \mathbb{R}^d , let $\Gamma_{\mu,\nu}$ denote the set of *couplings* between μ and ν so that $\gamma \in \Gamma_{\mu,\nu}$ iff $\gamma(U \times \mathbb{R}^d) = \mu(U)$ and $\gamma(\mathbb{R}^d \times V) = \nu(V)$.

For any $p \geq 1$, the *p-Wasserstein distance* W_p between μ and ν is defined as

$$(1) \quad W_p(\mu, \nu) := \inf_{\gamma \in \Gamma_{\mu,\nu}} \left(\int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^p d\gamma(x, y) \right)^{1/p}.$$

The definition of Wasserstein distances may be extended to measures defined on general metric spaces but such extensions are beyond the scope of this paper. We refer the reader to Villani [89] for a comprehensive treatment.

2.2 Spiked transport model

We introduce a new model that induces a low-dimensional structure on the optimal transport between two measures $\mu^{(1)}$ and $\mu^{(2)}$ over \mathbb{R}^d . To that end, fix a subspace $\mathcal{U} \subseteq \mathbb{R}^d$ of dimension $k \ll d$ and let $X^{(1)}, X^{(2)} \in \mathcal{U}$ be two random variables with arbitrary distributions. Next, let Z be a third random variable, independent of $(X^{(1)}, X^{(2)})$ and such that Z is supported on the orthogonal complement \mathcal{U}^\perp of \mathcal{U} . Finally, let

$$(2) \quad \begin{aligned} \mu^{(1)} &:= \text{Law}(X^{(1)} + Z) \\ \mu^{(2)} &:= \text{Law}(X^{(2)} + Z). \end{aligned}$$

Though $\mu^{(1)}$ and $\mu^{(2)}$ are high-dimensional distributions, they differ only on the low-dimensional subspace \mathcal{U} . Borrowing terminology from principal component analysis [47], we say that the pair $(\mu^{(1)}, \mu^{(2)})$ satisfies the *spiked transport model* and we call \mathcal{U} the *spike*.

We pose the following question: given n independent observations from both $\mu^{(1)}$ and $\mu^{(2)}$, is it possible to estimate the Wasserstein distance between them at a rate faster than $n^{-1/d}$?

2.3 Concentration assumptions

In order to establish sharp statistical results for estimation of the Wasserstein distance, it is necessary to adopt smoothness and decay assumptions on the measures in question [see, e.g., 9, 34]. We focus on a family of such conditions known as *transport inequalities*, the study of which is a central object in the theory of concentration of measure [56].

A probability measure μ on \mathbb{R}^d is said to satisfy the $T_p(\sigma^2)$ transport inequality if

$$(3) \quad W_p(\nu, \mu) \leq \sqrt{2\sigma^2 D(\nu \parallel \mu)}$$

for all probability measures ν on \mathbb{R}^d .

These inequalities interpolate between several well known assumptions in high-dimensional probability. For example inequality $T_1(\sigma^2)$ is essentially equivalent to the assertion that μ is subgaussian, and $T_2(\sigma^2)$ is implied by (and often equivalent to) a stronger log-Sobolev inequality [40].

Our main results on the estimation of $W_p(\mu^{(1)}, \mu^{(2)})$ are established under the assumption that both $\mu^{(1)}$ and $\mu^{(2)}$ satisfy a $T_p(\sigma^2)$ transport inequality. We show in Section 5 that (3) is precisely equivalent to requiring that the random variable $W_p(\mu_n, \mu)$ is subgaussian.

2.4 Wasserstein projection pursuit

To take advantage of the spiked transport model, we employ a natural estimation method that we call Wasserstein Projection Pursuit (WPP).

Let μ and ν be two probability distributions on \mathbb{R}^d . Given a $k \times d$ matrix U with orthonormal rows, let μ_U (resp. ν_U) denote the distribution of UY where $Y \sim \mu$ (resp. $Y \sim \nu$). We define

$$\tilde{W}_{p,k}(\mu, \nu) := \max_{U \in \mathcal{V}_k(\mathbb{R}^d)} W_p(\mu_U, \nu_U),$$

where the maximization is taken over the *Stiefel manifold* $\mathcal{V}_k(\mathbb{R}^d)$ of $k \times d$ matrices with orthonormal rows.

Given empirical measures $\mu_n^{(1)}$ and $\mu_n^{(2)}$ associated to $\mu^{(1)}$ and $\mu^{(2)}$ that satisfy the spiked transport model, we propose the following WPP estimator of $W_p(\mu^{(1)}, \mu^{(2)})$:

$$\hat{W}_{p,k} = \tilde{W}_{p,k}(\mu_n^{(1)}, \mu_n^{(2)}),$$

In the next section, we show that this estimator is near-minimax-optimal.

3. MAIN RESULTS

As a theoretical justification for Wasserstein projection pursuit, we prove that our procedure successfully avoids the curse of dimensionality under the spiked transport model. Our results primarily focus on the estimation of the Wasserstein distance itself but we also obtain as a byproduct of Wasserstein projection pursuit an estimator for the spike \mathcal{U} using standard perturbation results.

3.1 Estimation of the Wasserstein distance

The following theorem shows that Wasserstein projection pursuit takes advantage of the low-dimensional structure of the spiked transport model when estimating the Wasserstein distance.

THEOREM 1. *Let $(\mu^{(1)}, \mu^{(2)})$ satisfy the spiked transport model (2). For any $p \in [1, 2]$, if $\mu^{(1)}$ and $\mu^{(2)}$ satisfy the $T_p(\sigma^2)$ transport inequality, then the WPP estimator $\hat{W}_{p,k}$ satisfies*

$$\mathbb{E}|\hat{W}_{p,k} - W_p(\mu^{(1)}, \mu^{(2)})| \leq c_k \cdot \sigma \left(r_{p,k}(n) + \sqrt{\frac{d \log n}{n}} \right).$$

Strikingly, the rate $r_{p,d}(n)$ achieved by the naïve plug-in estimator (see Proposition 1) has been replaced by $r_{p,k}(n)$ —in other words, this estimator enjoys the rate typical for k -dimensional rather than d -dimensional measures. The only dependence on the ambient dimension is in the second term, which is of lower order than the first whenever $p > 1$ or $k > 2$. A more general version of this theorem appears in Section 6.

3.2 Estimation of the spike

We show that if the distance between $\mu^{(1)}$ and $\mu^{(2)}$ is large enough, Wasserstein projection pursuit recovers the subspace \mathcal{U} . For simplicity, we state here the result when $k = 1$ and defer the full version to Section 6.

THEOREM 2. *Let $(\mu^{(1)}, \mu^{(2)})$ satisfy the spiked transport model with $k = 1$ and let \mathcal{U} be spanned by the unit vector $u \in \mathbb{R}^d$. Fix $p \in [1, 2]$ and assume that $\mu^{(1)}$ and $\mu^{(2)}$ satisfy the $T_p(\sigma^2)$ transport inequality and that $W_p(\mu^{(1)}, \mu^{(2)}) \gtrsim 1$. Then the estimator*

$$\hat{u} := \operatorname{argmax}_{v \in \mathbb{R}^d, \|v\|=1} W_p(\mu_v^{(1)}, \mu_v^{(2)})$$

satisfies

$$\mathbb{E} \sin^2(\angle(\hat{u}, u)) \lesssim \sigma \cdot \left(n^{-1/2p} + \sqrt{\frac{d \log n}{n}} \right).$$

3.3 Lower bounds

To show that Theorem 1 has the right dependence on n and d , we exhibit two lower bounds, which imply that neither term in Theorem 1 can be avoided.

To show the optimality of the first term, we define

$$r'_{p,d}(n) := c_{p,d} \begin{cases} n^{-1/2p} & \text{if } d < 2p \\ n^{-1/2p} & \text{if } d = 2p \\ (n \log n)^{-1/d} & \text{if } d > 2p \end{cases}$$

THEOREM 3. *Fix $p \geq 1$. For any estimator \hat{W} , there exists a pair of measures μ and ν supported on $[0, 1]^d$ such*

$$\mathbb{E}|\hat{W} - W_p(\mu^{(1)}, \mu^{(2)})| \geq r'_{p,d}(n).$$

This lower bound readily implies that the plug-in estimator for the Wasserstein distance is optimal up to logarithmic factors. By embedding $[0, 1]^k$ into $[0, 1]^d$, this result likewise implies that the term $r_{p,k}(n)$ in Theorem 1 is essentially optimal. A proof appears in Section 7

Independently, Liang [58] recently obtained a similar result in the case $p = 1$. More specifically, he proved that when $d \geq 2$, for any estimator \hat{W} , there exist probability measures $\mu^{(1)}$ and $\mu^{(2)}$ such that the following lower bound holds:

$$\mathbb{E}|\hat{W} - W_1(\mu^{(1)}, \mu^{(2)})| \gtrsim \frac{\log \log n}{\log n} n^{-1/d}$$

In particular, while our lower bound is slightly stronger and holds for all $p \geq 1$, both our result and that of Liang [58] fail to match the naïve upper bound of order $n^{1/d}$ by logarithmic factors when d is large. The presence of a logarithmic factor in our lower bound comes from a reduction to estimating the total variation distance. In that case, as in several other instances of functional estimation problems, the presence of this factor is, in fact, optimal, and has been dubbed *sample size enlargement* [45]. Closing this gap in the context of estimation of the Wasserstein distance is an interesting and fundamental question.

The only appearance of the ambient dimension d is in the second term of Theorem 1. The following theorem shows that this dependence cannot be eliminated, even when $k = 1$.

THEOREM 4. *Let $p \in [1, 2]$ and $\sigma > 0$, and assume $k = 1$. For all estimators \hat{W} , there exists a pair of measures $\mu^{(1)}$ and $\mu^{(2)}$ satisfying the spiked transport model and $T_p(\sigma^2)$ such that*

$$\mathbb{E}|\hat{W} - W_p(\mu^{(1)}, \mu^{(2)})| \gtrsim \sigma \sqrt{\frac{d}{n}}.$$

The proof of Theorem 4 is deferred to the appendix. For problems where d is large, dependence on d may be a crippling limitation. In that case, we conjecture that assuming a sparse spike, in the same spirit as sparse PCA, can mitigate this effect and bring interpretability to the estimated spike.

3.4 A computational-statistical gap

The WPP estimator achieving the rate in Theorem 1 is computationally expensive to implement, which raises the question of whether an efficient estimator exists achieving the same rate. We give evidence in the form of a statistical query lower bound that no such estimator exists. The statistical query model considers algorithms with access to an oracle $VSTAT(t)$, where $t > 0$ is a parameter which plays the role of sample size. We show that any such algorithm for estimating the Wasserstein distance needs an *exponential* number of queries to an oracle with *exponential* sample size parameter, even under the spiked transport model. By contrast, Theorem 1 implies that a non-efficient estimator needs a number of samples only polynomial in the dimension.

THEOREM 5. *Let $p \in [1, 2]$, and consider probability measures $\mu^{(1)}$ and $\mu^{(2)}$ satisfying the spiked transport model. There exists a positive constant c such that any statistical query algorithm which estimates $W_p(\mu^{(1)}, \mu^{(2)})$ to accuracy $1/\text{poly}(d)$ with probability at least $2/3$ requires at least 2^{cd} queries to $VSTAT(2^{cd})$.*

4. LOW-DIMENSIONAL WASSERSTEIN DISTANCES

Motivated by projection pursuit, we define the following version of the Wasserstein distance which measures the discrepancy between low-dimensional projections of the measures.

DEFINITION 1. For $k \in [d]$, the *k-dimensional Wasserstein distance* between $\mu^{(1)}$ and $\mu^{(2)}$ is

$$\tilde{W}_{p,k}(\mu^{(1)}, \mu^{(2)}) := \sup_{U \in \mathcal{V}_k(\mathbb{R}^d)} W_p(\mu_U^{(1)}, \mu_U^{(2)}).$$

This definition has been proposed independently and concurrently by a number of other recent works [24, 50, 66].

We will use throughout the following basic fact about the k -dimensional Wasserstein distance [66, Proposition 1].

PROPOSITION 2. *$\tilde{W}_{p,k}$ is a metric on the set of probability measures over \mathbb{R}^d with finite p th moment.*

The definition of $\tilde{W}_{p,k}$ is chosen so that, under the spiked transport model (2), the k -dimensional Wasserstein distance agrees with the normal Wasserstein distance.

PROPOSITION 3. *Under the spiked transport model (2),*

$$\tilde{W}_{p,k}(\mu^{(1)}, \mu^{(2)}) = W_p(\mu^{(1)}, \mu^{(2)}).$$

Proposition 3 follows from the following statement, which pertains to distributions that are allowed to have a different component on the space orthogonal to the spike \mathcal{U} . Suppose $\nu^{(1)}$ and $\nu^{(2)}$ satisfy

$$(4) \quad \begin{aligned} \nu^{(1)} &= \text{Law}(X^{(1)} + Z^{(1)}) \\ \nu^{(2)} &= \text{Law}(X^{(2)} + Z^{(2)}), \end{aligned}$$

where as before $X^{(1)}$ and $X^{(2)}$ are supported on a subspace \mathcal{U} and $Z^{(1)}$ and $Z^{(2)}$ are supported on its orthogonal complement \mathcal{U}^\perp , and where we assume that $X^{(i)}$ and $Z^{(i)}$ are independent for $i \in \{1, 2\}$. Note that unlike in the spiked transport model, the components $Z^{(1)}$ and $Z^{(2)}$ on the orthogonal complement of \mathcal{U} need not be identical. The following result shows that under this relaxed model, the k -dimensional Wasserstein distance between $\nu^{(1)}$ and $\nu^{(2)}$ still captures the true Wasserstein distance between the distributions as long as the distributions of $Z^{(1)}$ and $Z^{(2)}$ are sufficiently close.

PROPOSITION 4. *Under the relaxed spiked transport model (4),*

$$|\tilde{W}_{p,k}(\nu^{(1)}, \nu^{(2)}) - W_p(\nu^{(1)}, \nu^{(2)})| \leq W_p(\text{Law}(Z^{(1)}), \text{Law}(Z^{(2)})).$$

5. CONCENTRATION

A key step to establish the upper bound of Section 6 consists in establishing good concentration properties for the Wasserstein distance between a measure and its empirical counterpart. The main assumption we adopt is that the measures in question satisfy a *transport inequality*. Since the pioneering work of [61, 62] and [81], transport inequalities have played a central role in the analysis of the concentration properties of high-dimensional measures.

We require two definitions.

DEFINITION 2. Given a Polish space \mathcal{X} equipped with a metric ρ , denote by $\mathcal{P}(\mathcal{X})$ the space of all Borel probability measures \mathcal{X} . Let $\mathcal{P}_p(\mathcal{X}) := \{\mu \in \mathcal{P}(\mathcal{X}) : \int \rho(x, \cdot)^p d\mu(x) < \infty\}$.

A measure $\mu \in \mathcal{P}_p(\mathcal{X})$ satisfies the $T_p(\sigma^2)$ inequality for some $\sigma > 0$ if

$$W_p(\nu, \mu) \leq \sqrt{2\sigma^2 D(\nu \parallel \mu)} \quad \forall \nu \in \mathcal{P}_p(\mathcal{X}),$$

where W_p is the Wasserstein- p distance on (\mathcal{X}, ρ) and D is the Kullback-Leibler divergence.

DEFINITION 3. A random variable X on \mathbb{R} is σ^2 -subgaussian if

$$\mathbb{E}e^{\lambda(X - \mathbb{E}X)} \leq e^{\lambda^2 \sigma^2 / 2} \quad \forall \lambda \in \mathbb{R}.$$

In this section, we present a surprisingly simple equivalence between transport inequalities and subgaussian concentration for the Wasserstein distance. The essence of this result is present in the works of [Gozlan and Léonard](#) [see 38, 39, 40], and similar bounds have been obtained by Bolley et al. [12]. Nevertheless, we could not find this simple fact stated in a form suitable for our purposes in the literature. For any measure μ , recall that the random measure $\mu_n := \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$, where $X_i \sim \mu$ i.i.d., denotes its associated empirical measure.

THEOREM 6. *Let $p \in [1, 2]$. A measure $\mu \in \mathcal{P}_p(\mathcal{X})$ satisfies $T_p(\sigma^2)$ if and only if the random variable $W_p(\mu_n, \mu)$ is σ^2/n -subgaussian for all n .*

Because $T_{p'}(\sigma^2)$ implies $T_p(\sigma^2)$ when $p' \geq p$, Theorem 6 also implies that a measure satisfying $T_{p'}(\sigma^2)$ has good concentration for W_p if $p \leq p'$. In the opposite direction, if $p > p'$, then satisfying $T_{p'}(\sigma^2)$ still yields a weaker concentration bound. A modification of the proof of Theorem 6 yields the following result.

THEOREM 7. *Let $p' \in [1, 2]$ and $p \geq 1$. If $\mu \in \mathcal{P}_{p'}(\mathcal{X})$ satisfies $T_{p'}(\sigma^2)$, then $W_p(\mu_n, \mu)$ is $\sigma^2/n^{1 - \left(\frac{2}{p'} - \frac{2}{p}\right)_+}$ subgaussian.*

The conclusion of Theorem 7 is interesting whenever $\left(\frac{2}{p'} - \frac{2}{p}\right)_+ < 1$. For example, if we assume merely that μ satisfies $T_1(\sigma^2)$, Theorem 7 only yields a nontrivial concentration result for $W_p(\mu_n, \mu)$ when $p < 2$; by contrast, if μ satisfies $T_2(\sigma^2)$, then Theorem 7 implies a concentration result for $W_p(\mu_n, \mu)$ for all $p < \infty$.

In Section 6, we require concentration properties not of the Wasserstein distance itself but of the k -dimensional Wasserstein distance. The following result shows that low-dimensional projections inherit the concentration properties of the d -dimensional measure.

PROPOSITION 5. *Let $U \in \mathcal{V}_k(\mathbb{R}^d)$. For any $p \in [1, 2]$ and $\sigma > 0$, if μ satisfies $T_p(\sigma^2)$, then so does μ_U .*

PROOF. The projection $x \mapsto Ux$ is a contraction. The result then follows from Gozlan [37, Corollary 20] [see also 63]. \square

We conclude this section by giving some simple conditions under which the $T_1(\sigma^2)$ inequality is satisfied. The following characterization is well known. Denote by $\text{Lip}(\mathcal{X})$ the space of all functions $f : \mathcal{X} \rightarrow \mathbb{R}$ satisfying $|f(x) - f(y)| \leq d(x, y)$ for all $x, y \in \mathbb{R}$.

PROPOSITION 6 (10, Theorem 1.3). *A measure $\mu \in \mathcal{P}_1(\mathcal{X})$ satisfies $T_1(\sigma^2)$ if and only if $f(X)$ is σ^2 -subgaussian for all $f \in \text{Lip}(\mathcal{X})$.*

It is common to extend Definition 3 to random vectors as follows.

DEFINITION 4. A random variable X on \mathbb{R}^d is σ^2 -subgaussian if $u^\top X$ is σ^2 -subgaussian for all $u \in \mathbb{R}^d$ satisfying $\|u\| = 1$.

Subgaussian random vectors yield a large collection of random variables satisfying a T_1 inequality.

LEMMA 1. *If μ on \mathbb{R}^k satisfies $T_p(\sigma^2)$ for any $p \geq 1$, then $X \sim \mu$ is σ^2 -subgaussian. Conversely, if $X \sim \mu$ on \mathbb{R}^k is σ^2 -subgaussian, then μ satisfies $T_1(Ck\sigma^2)$ for a universal constant $C > 0$.*

If the entries of X are independent, then the result holds with $C = 1$ by a result of Marton [61]. The presence of the factor k in the converse statement is unavoidable; unlike T_2 inequalities, T_1 inequalities do not exhibit dimension-free concentration [40].

6. UPPER BOUNDS

In this section, we establish that under the spiked transport model, Wasserstein projection pursuit produces a significantly more accurate estimate of the Wasserstein distance than the plug-in estimator.

Let $\mu^{(1)}$ and $\mu^{(2)}$ be two measures generated according to the spiked transport model (2). For $i \in \{1, 2\}$, we let $\mu_n^{(i)} := \frac{1}{n} \sum_{j=1}^n \delta_{X_j^{(i)}}$, where $X_j^{(i)} \sim \mu^{(i)}$ are i.i.d. We define

$$\hat{W}_{p,k} := \tilde{W}_{p,k}(\mu_n^{(1)}, \mu_n^{(2)}).$$

Our main upper bound shows that $\hat{W}_{p,k}$ converges to the true Wasserstein distance $W_p(\mu, \nu)$ at a rate much faster than $n^{-1/d}$.

THEOREM 8. *Let $p' \in [1, 2]$ and $p \geq 1$. Under the spiked transport model, if $\mu^{(1)}$ and $\mu^{(2)}$ satisfy $T_{p'}(\sigma^2)$, then*

$$\mathbb{E}|\hat{W}_{p,k} - W_p(\mu^{(1)}, \mu^{(2)})| \lesssim \sigma \left(r_{p,k}(n) + c_p \cdot n^{(\frac{1}{p'} - \frac{1}{p})_+} \sqrt{\frac{dk \log n}{n}} \right).$$

Theorem 8 can also be extended to the misspecified model proposed in (4)

THEOREM 9. *Let $p' \in [1, 2]$ and $p \geq 1$. Under the relaxed spiked transport model (4), if $\nu^{(1)}$ and $\nu^{(2)}$ satisfy $T_{p'}(\sigma^2)$, then*

$$\mathbb{E}|\hat{W}_{p,k} - W_p(\nu^{(1)}, \nu^{(2)})| \lesssim \sigma \left(r_{p,k}(n) + c_p \cdot n^{(\frac{1}{p'} - \frac{1}{p})_+} \sqrt{\frac{dk \log n}{n}} \right) + \varepsilon,$$

where $\varepsilon = W_p(\text{Law}(Z^{(1)}), \text{Law}(Z^{(2)}))$.

Theorem 9, which follows almost immediately from the proof of Theorem 8, establishes that Wasserstein projection pursuit brings statistical benefits even in the situation where the spiked transport model holds only approximately.

Theorem 8 follows from the following two propositions. We first show that the quality of the proposed estimator $\hat{W}_{p,k}$ can be bounded by the sum of two terms depending only on $\mu_n^{(1)}$ and $\mu_n^{(2)}$ individually.

PROPOSITION 7.

$$\mathbb{E}|\hat{W}_p - W_p(\mu^{(1)}, \mu^{(2)})| \leq \mathbb{E}\tilde{W}_{p,k}(\mu^{(1)}, \mu_n^{(1)}) + \mathbb{E}\tilde{W}_{p,k}(\mu^{(2)}, \mu_n^{(2)})$$

PROOF. Since $\hat{W}_p = \tilde{W}_{p,k}(\mu_n^{(1)}, \mu_n^{(2)})$ and $W_p(\mu^{(1)}, \mu^{(2)}) = \tilde{W}_{p,k}(\mu^{(1)}, \mu^{(2)})$, the claim is immediate from Proposition 2. \square

The following proposition allows us to bound both terms of Proposition 7 by the desired quantity.

PROPOSITION 8. *Let $p' \in [1, 2]$ and $p \geq 1$. If μ satisfies $T_{p'}(\sigma^2)$, then*

$$\mathbb{E}\tilde{W}_{p,k}(\mu, \mu_n) \lesssim \sigma \left(r_{p,k}(n) + c_p \cdot n^{(\frac{1}{p'} - \frac{1}{p})_+} \sqrt{\frac{dk \log n}{n}} \right).$$

PROOF. Wasserstein distances are invariant under translating both measures by the same vector. Therefore, we can assume without loss of generality that μ has mean 0. Likewise, by homogeneity, we assume $\sigma = 1$.

Let $Z_U := W_p(\mu_U, (\mu_n)_U)$. We first show that the process Z_U is Lipschitz.

LEMMA 2. *There exists a random variable L such that for all $U, V \in \mathcal{V}_k(\mathbb{R}^d)$,*

$$|Z_U - Z_V| \leq L \|U - V\|_{\text{op}}$$

and $\mathbb{E}L \lesssim \sqrt{dkp}$.

PROOF. Let $X \sim \mu$. Then

$$\begin{aligned} |Z_U - Z_V| &\leq W_p(\mu_U, \mu_V) + W_p((\mu_n)_U, (\mu_n)_V) \\ &\leq (\mathbb{E}\|(U - V)X\|^p)^{1/p} + \left(\frac{1}{n} \sum_{i=1}^n \|(U - V)X_i\|^p\right)^{1/p} \\ &\leq \|U - V\|_{\text{op}} \left((\mathbb{E}\|X\|^p)^{1/p} + \left(\frac{1}{n} \sum_{i=1}^n \|X_i\|^p\right)^{1/p} \right). \end{aligned}$$

We obtain that

$$|Z_U - Z_V| \leq L \|U - V\|_{\text{op}}$$

where $L = (\mathbb{E}\|X\|^p)^{1/p} + \left(\frac{1}{n} \sum_{i=1}^n \|X_i\|^p\right)^{1/p}$. By Jensen's inequality, we have $\mathbb{E}L \leq 2(\mathbb{E}\|X\|^p)^{1/p}$. Together with Lemma 5, it yields the claim. \square

By Theorem 7, for all $U \in \mathcal{V}_k(\mathbb{R}^d)$, the random variable Z_U is $n^{-1 + \left(\frac{2}{p'} - \frac{2}{p}\right)_+}$ sub-gaussian. Therefore, by a standard ε -net argument, if we denote by $\mathcal{N}(\mathcal{V}_k, \varepsilon, \|\cdot\|_{\text{op}})$ the covering number of \mathcal{V}_k with respect to the operator norm, we obtain

$$\mathbb{E} \sup_{U \in \mathcal{V}_k(\mathbb{R}^d)} (Z_U - \mathbb{E}Z_U) \lesssim \inf_{\varepsilon > 0} \left\{ \varepsilon \mathbb{E}L + n^{\left(\frac{1}{p'} - \frac{1}{p}\right)_+} \sqrt{\frac{\log \mathcal{N}(\mathcal{V}_k, \varepsilon, \|\cdot\|_{\text{op}})}{n}} \right\}.$$

Lemma 4 shows that there exists a universal constant c such that $\log \mathcal{N}(\mathcal{V}_k, \varepsilon, \|\cdot\|_{\text{op}}) \leq dk \log \frac{c\sqrt{k}}{\varepsilon}$ for $\varepsilon \in (0, 1]$. Choosing $\varepsilon = \sqrt{k/n}$ yields

$$\begin{aligned} \mathbb{E} \sup_{U \in \mathcal{V}_k(\mathbb{R}^d)} (Z_U - \mathbb{E}Z_U) &\lesssim \sqrt{\frac{dkp}{n}} + n^{\left(\frac{1}{p'} - \frac{1}{p}\right)_+} \sqrt{\frac{dk \log n}{n}} \\ &\leq c_p \cdot n^{\left(\frac{1}{p'} - \frac{1}{p}\right)_+} \sqrt{\frac{dk \log n}{n}}. \end{aligned}$$

Applying Proposition 17 yields

$$\begin{aligned} \mathbb{E} \sup_{U \in \mathcal{V}_k(\mathbb{R}^d)} W_p(\mu_U, (\mu_n)_U) &\leq \sup_{U \in \mathcal{V}_k(\mathbb{R}^d)} \mathbb{E}W_p(\mu_U, (\mu_n)_U) + \mathbb{E} \sup_{U \in \mathcal{V}_k(\mathbb{R}^d)} (Z_U - \mathbb{E}Z_U) \\ &\lesssim r_{p,k}(n) + c_p \cdot n^{\left(\frac{1}{p'} - \frac{1}{p}\right)_+} \sqrt{\frac{dk \log n}{n}}, \end{aligned}$$

as claimed. \square

We also obtain a Davis-Kahan-type theorem on subspace recovery. Given two subspaces \mathcal{U}_1 and \mathcal{U}_2 , the *minimal angle* [25, 27] between them is defined to be

$$\angle(\mathcal{U}_1, \mathcal{U}_2) := \arccos \left(\sup_{u_1 \in \mathcal{U}_1, u_2 \in \mathcal{U}_2} \frac{u_1^\top u_2}{\|u_1\| \|u_2\|} \right).$$

If $\angle(\mathcal{U}_1, \mathcal{U}_2) = 0$, then $\mathcal{U}_1 \cap \mathcal{U}_2 \neq \{0\}$, so that \mathcal{U}_1 and \mathcal{U}_2 are at least partially aligned. In the important special case that \mathcal{U}_1 and \mathcal{U}_2 are each one dimensional, this definition reduces to the angle between the subspaces.

The following result indicates that as long as $\mu^{(1)}$ and $\mu^{(2)}$ are well separated, Wasserstein projection pursuit also yields a subspace with at least partial alignment to \mathcal{U} .

THEOREM 10. *Let $p' \in [1, 2]$ and $p \geq 1$. Assume that $\mu^{(1)}$ and $\mu^{(2)}$ satisfy the spiked transport model and $T_{p'}(\sigma^2)$. Let $\hat{\mathcal{U}} := \text{span}(\hat{U})$, where*

$$\hat{U} := \operatorname{argmax}_{U \in \mathcal{V}_k(\mathbb{R}^d)} W_p((\mu_n^{(1)})_U, (\mu_n^{(2)})_U).$$

Then

$$\mathbb{E} \sin^2(\angle(\hat{\mathcal{U}}, \mathcal{U})) \lesssim \frac{\sigma \left(r_{p,k}(n) + c_p \cdot n^{(\frac{1}{p'} - \frac{1}{p})_+} \sqrt{\frac{dk \log n}{n}} \right)}{W_p(\mu^{(1)}, \mu^{(2)})}.$$

A proof of Theorem 10 appears in the appendix. Note that the a bound on the minimal angle is a rather weak guarantee. Indeed, $\angle(\hat{\mathcal{U}}, \mathcal{V}) \rightarrow 0$ implies that the subspaces $\hat{\mathcal{U}}$ and \mathcal{U} share at least a common line asymptotically but not more. When $k = 1$, this ensures recovery of the subspace, but this no longer holds true for higher dimensional spikes. In retrospect such a guarantee is all that we can hope for under the mere assumption that $W_p(\mu^{(1)}, \mu^{(2)}) > 0$. Indeed, it may be the case that these distributions differ only on a one dimensional space. Stronger guarantees may be achieved by assuming that $W_p(\mu_V^{(1)}, \mu_V^{(2)}) > 0$ for a large family of V , but we leave them for future research.

7. A LOWER BOUND ON ESTIMATING THE WASSERSTEIN DISTANCE

In this section, we prove that the rate $r_{p,d}(n)$ is optimal for estimating the Wasserstein distance, up to logarithmic factors. The core idea of our lower bound is to relate estimating the Wasserstein distance to the problem of estimating total variation distance, sharp rates for which are known [45, 86]. To obtain sufficient control over the Wasserstein distance as a function of total variation, we prove a refined bound incorporating both total variation and the χ^2 divergence (Proposition 9). We then show a modified lower bound (Proposition 10) for a testing problem involving the total variation distance over the class of distributions on $[m]$ close to the uniform measure in χ^2 divergence.

In the interest of generality, we formulate our results for any compact metric space \mathcal{X} whose covering numbers satisfy

$$(5) \quad c\varepsilon^{-d} \leq \mathcal{N}(\mathcal{X}, \varepsilon) \leq C\varepsilon^{-d}$$

for all $\varepsilon \leq \text{diam}(\mathcal{X})$. This condition clearly holds for compact subsets of \mathbb{R}^d and more generally for metric spaces with Minkowski dimension d . We adopt the assumption $\text{diam}(\mathcal{X}) = 1$ without loss of generality.

Let \mathcal{P} be the set of distributions supported on \mathcal{X} and let $R(n, \mathcal{P})$ denote the minimax risk over \mathcal{P} ,

$$R(n, \mathcal{P}) := \inf_W \sup_{\mu, \nu \in \mathcal{P}} \mathbb{E}_{\mu, \nu} |\hat{W} - W_p(\mu, \nu)|.$$

The bound $R(n, \mathcal{P}) \gtrsim n^{-1/2p}$ is an almost trivial consequence of the fact that the distribution $\frac{1}{2}\delta_{-1} + \frac{1}{2}\delta_1$ cannot be distinguished from $(\frac{1}{2} + \varepsilon)\delta_{-1} + (\frac{1}{2} - \varepsilon)\delta_1$ on the basis of n samples when $\varepsilon \asymp n^{-1/2}$. The interesting part of Theorem 3 is the rate when $d > 2p$. We prove the following.

THEOREM 11. *Let $d > 2p \geq 2$ and assume \mathcal{X} satisfies (5). Then*

$$R(n, \mathcal{P}) \geq C_{d,p}(n \log n)^{-1/d}.$$

Before proving Theorem 11, we establish the two propositions described above. Proposition 9 allows us to reduce Theorem 11 to an estimation problem involving total variation distance, and Proposition 10 is a lower bound on the minimax rate for that total variation estimation problem.

PROPOSITION 9. *Assume $d > 2p \geq 2$, and let m be a positive integer. Let u be the uniform distribution on $[m] := \{1, \dots, m\}$. There exists a random function $F : [m] \rightarrow X$ such that for any distribution q on $[m]$,*

$$cm^{-1/d} d_{\text{TV}}(q, u)^{\frac{1}{p}} \leq W_p(F_{\#}q, F_{\#}u) \leq C_{d,p} m^{-1/d} (\chi^2(q, u))^{1/d} d_{\text{TV}}(q, u)^{\frac{1}{p} - \frac{2}{d}}$$

with probability at least .9.

PROOF. Lemma 6 shows that the condition $\mathcal{N}(\mathcal{X}, \varepsilon) \geq c\varepsilon^{-d}$ implies the existence a set $\mathcal{G}_m := \{x_1, \dots, x_m\} \subseteq \mathcal{X}$ such that $d(x_i, x_j) \gtrsim m^{-1/d}$ for all $i \neq j$. We select F uniformly at random from the set of all bijections from $[m]$ to \mathcal{G}_m .

To show the lower bound, we note that any points $x, y \in \mathcal{G}_m$ satisfy

$$d(x, y)^p \gtrsim m^{-p/d} \mathbb{1}\{x \neq y\},$$

which implies that for any coupling π between $F_{\#}q$ and $F_{\#}u$

$$\begin{aligned} \int d(x, y)^p d\pi(x, y) &\gtrsim m^{-p/d} \mathbb{P}_{\pi}[X \neq Y] \\ &\geq m^{-p/d} d_{\text{TV}}(F_{\#}q, F_{\#}u) \\ &= m^{-p/d} d_{\text{TV}}(q, u). \end{aligned}$$

The lower bound therefore holds with probability 1.

We now turn to the upper bound. We employ a dyadic covering bound [90, Proposition 1]. For any k^* there exists a dyadic partition $\{Q^k\}_{1 \leq k \leq k^*}$ of X with parameter $\delta = 1/3$ such that $|Q^k| \leq \mathcal{N}(\mathcal{X}, 3^{-(k+1)})$. We obtain that for any $k^* \geq 0$,

$$W_p^p(F_{\#}q, F_{\#}u) \leq 3^{-k^*p} + \sum_{k=1}^{k^*} 3^{-(k-1)p} \sum_{Q_i^k \in \mathcal{Q}^k} |F_{\#}q(Q_i^k) - F_{\#}u(Q_i^k)|,$$

By Lemma 7, for any k ,

$$\mathbb{E} \sum_{Q_i^k \in \mathcal{Q}^k} |F_{\#}q(Q_i^k) - F_{\#}u(Q_i^k)| \leq 2d_{\text{TV}}(q, u) \wedge C_{d,p} \left(\frac{3^{kd} \chi^2(q, u)}{m} \right)^{1/2}.$$

Let k_0 be a positive integer to be fixed later.

By applying the first bound, we obtain

$$\mathbb{E} \sum_{k > k_0} 3^{-(k-1)p} \sum_{Q_i^k \in \mathcal{Q}^k} |F_{\#}q(Q_i^k) - F_{\#}u(Q_i^k)| \lesssim 3^{-k_0p} d_{\text{TV}}(q, u).$$

Applying the second bound and recalling that $d/2 > p$ yields

$$\begin{aligned} \mathbb{E} \sum_{k \leq k_0} 3^{-(k-1)p} \sum_{Q_i^k \in \mathcal{Q}^k} |F_{\#}q(Q_i^k) - F_{\#}u(Q_i^k)| &\leq C_{d,p} \left(\frac{\chi^2(q, u)}{m} \right)^{1/2} \sum_{k \leq k_0} 3^{k(d/2-p)} \\ &\leq C_{d,p} \left(\frac{\chi^2(q, u)}{m} \right)^{1/2} 3^{k_0(d/2-p)}, \end{aligned}$$

We obtain for any $k^* \geq 0$ that

$$\mathbb{E} W_p^p(F_{\#}q, F_{\#}u) \leq C_{d,p} \left(\frac{\chi^2(q, u)}{m} \right)^{1/2} 3^{k_0(d/2-p)} + C \cdot 3^{-k_0p} d_{\text{TV}}(q, u) + 3^{-k^*p},$$

and taking $k^* \rightarrow \infty$ it suffices to bound the first two terms.

Let k_0 to be the smallest positive integer such that

$$3^{k_0d} \geq m \frac{d_{\text{TV}}(q, u)^2}{\chi^2(q, u)}.$$

Then

$$3^{k_0d/2} \left(\frac{\chi^2(q, u)}{m} \right)^{1/2} \leq C_d \cdot d_{\text{TV}}(q, u),$$

and hence

$$\mathbb{E} W_p^p(F_{\#}q, F_{\#}u) \leq C_{d,p} \cdot 3^{-k_0p} d_{\text{TV}}(q, u) \leq C_{d,p} m^{-p/d} (\chi^2(q, u))^{p/d} d_{\text{TV}}(q, u)^{1-\frac{2p}{d}}.$$

The claim follows from Markov's inequality. \square

We now show that there are composite hypotheses that are well separated in total variation distance but nevertheless hard to distinguish on the basis of samples.

PROPOSITION 10. *Fix a positive integer n and a constant $\delta \in [0, 1/10]$. Given a positive integer m , let \mathcal{D}_m be the set of probability distributions q on $[m]$ satisfying $\chi^2(q, u) \leq 9$. Denote by $\mathcal{D}_{m,\delta}^-$ the subset of \mathcal{D}_m of distributions satisfying $d_{\text{TV}}(q, u) \leq \delta$ and by \mathcal{D}_m^+ the subset of \mathcal{D}_m satisfying $d_{\text{TV}}(q, u) \geq 1/4$. If $m = \lceil C\delta^{-1}n \log n \rceil$ for a sufficiently large universal constant C and n is sufficiently large, then*

$$\inf_{\psi} \left\{ \sup_{q \in \mathcal{D}_m^+} \mathbb{P}_q[\psi = 1] + \sup_{q \in \mathcal{D}_{m,\delta}^-} \mathbb{P}_q[\psi = 0] \right\} \geq .9,$$

where the infimum is taken over all (possibly randomized) tests based on n samples.

The proof of Proposition 10 follows a strategy due to Valiant and Valiant [85] and Wu and Yang [92], and our argument is a modification of theirs which permits simultaneous control of total variation and the χ^2 divergence. We give the proof in Appendix A.2.

We now give a proof of the main theorem.

PROOF OF THEOREM 11. Let \hat{W} be any estimator for the Wasserstein distance between distributions on \mathcal{X} constructed on the basis of n samples from each distribution.

Let u be the uniform distribution on $[m]$, for some m to be specified. Let c^* be the constant appearing in the lower bound of Proposition 9 and define $\Delta_d = \frac{1}{16}c^*m^{-1/d}$. Given n samples X_1, \dots, X_n from an unknown distribution on $[m]$, define the randomized test

$$\psi = \psi(X_1, \dots, X_n) := \mathbb{1}\{\hat{W}(F(X_1), \dots, F(X_n); F(Y_1), \dots, F(Y_n)) \leq 2\Delta_d\},$$

where F is the random function constructed in Proposition 9 and where Y_i are i.i.d. from u .

By Proposition 9, if $\delta \leq \delta_{d,p} := \left(\frac{c^*}{176C_{d,p}}\right)^{1/(p-2/d)}$, any $q \in \mathcal{D}_{m,\delta}^-$ satisfies the bound $W_p(F_{\#}q, F_{\#}u) \leq \Delta_d$ with probability at least .9. Likewise, for $q \in \mathcal{D}_m^+$, the bound $W_p(F_{\#}q, F_{\#}u) \geq 3\Delta_d$ also holds with probability at least .9.

Define the event $A = \{|\hat{W} - W_p(F_{\#}q, F_{\#}u)| \geq \Delta_d\}$. We obtain, for any $q \in \mathcal{D}_{m,\delta}^-$,

$$\begin{aligned} \mathbb{E}_F \mathbb{P}_{F_{\#}q, F_{\#}u}[A] &\geq \mathbb{E}_F \mathbb{P}_{F_{\#}q, F_{\#}u}[\hat{W} > 2\Delta_d \text{ and } W_p(F_{\#}q, F_{\#}u) \leq \Delta_d] \\ &\geq \mathbb{E}_F \mathbb{P}_{F_{\#}q, F_{\#}u}[\hat{W} > 2\Delta_d] - \mathbb{P}[W_p(F_{\#}q, F_{\#}u) > \Delta_d] \\ &\geq \mathbb{P}_q[\psi = 0] - .1, \end{aligned}$$

and analogously for $q \in \mathcal{D}_m^+$,

$$\mathbb{E}_F \mathbb{P}_{F_{\#}q, F_{\#}u}[A] \geq \mathbb{P}_q[\psi = 1] - .1.$$

For any estimator \hat{W} , we have

$$\begin{aligned} \sup_{\mu, \nu \in \mathcal{P}} \mathbb{P}_{\mu, \nu}[|\hat{W} - W_p(\mu, \nu)| \geq \Delta_d] &\geq \frac{1}{2} \left(\sup_{q \in \mathcal{D}_m^+} \mathbb{E}_F \mathbb{P}_{F_{\#}q, F_{\#}u}[A] + \sup_{q \in \mathcal{D}_{m,\delta}^-} \mathbb{E}_F \mathbb{P}_{F_{\#}q, F_{\#}u}[A] \right) \\ &\geq \frac{1}{2} \left(\sup_{q \in \mathcal{D}_m^+} \mathbb{P}_q[\psi = 1] + \sup_{q \in \mathcal{D}_{m,\delta}^-} \mathbb{P}_q[\psi = 0] \right) - .1. \end{aligned}$$

Choosing $m = \lceil C\delta^{-1}n \log n \rceil$ for a sufficiently large constant C and applying Proposition 10 yields that $\sup_{\mu, \nu \in \mathcal{P}} \mathbb{P}_{\mu, \nu}[|\hat{W} - W_p(F_{\#}q, F_{\#}u)| \geq \Delta_d] \geq .8$, and Markov's inequality yields the claim. \square

8. COMPUTATIONAL-STATISTICAL GAPS FOR THE SPIKED TRANSPORT MODEL

Sections 4 and 7 clarify the statistical price for estimating the Wasserstein distance for high-dimensional measures. Section 4 shows that the curse of dimensionality can be avoided under the spiked transport model. The WPP estimator exploits the low-dimensional structure in the spiked transport model, thereby

beating the worst-case rate presented in Section 7. However, it is not clear how to make the estimator we propose computationally efficient. In this section, we give evidence that this obstruction is a fundamental obstacle, that is, that no computationally efficient estimator can beat the curse of dimensionality.

The statistical query model, first introduced in the context of PAC learning [48], is a well known computational framework for analyzing statistical algorithms. Instead of being given access to data points from a distribution, a statistical query (SQ) algorithm can approximately evaluate the expectation of arbitrary functions with respect to the distribution. This model naturally captures the power of noise-tolerant algorithms [48] and is strong enough to implement nearly all common machine learning procedures [see, e.g. 8].

We recall the following definition.

DEFINITION 5. Given a distribution \mathcal{D} on \mathbb{R}^d , for any sample size parameter $t > 0$ and function $f : \mathbb{R}^d \rightarrow [0, 1]$, the oracle $\text{VSTAT}(t)$ returns a value $v \in [p - \tau, p + \tau]$, where $p = \mathbb{E}f(X)$ and $\tau = \frac{1}{t} \vee \sqrt{\frac{p(1-p)}{t}}$.

A query to a $\text{VSTAT}(t)$ oracle can be simulated by using a data set of size approximately t . Our main result proves a lower bound against an oracle with sample size parameter $t = 2^{cd}$ for a positive constant c . Simulating such an oracle would require a number of samples exponential in the dimension. Nevertheless, we show that even under this strong assumption, at least 2^{cd} queries to the oracle are required. This result suggests that any *computationally efficient* procedure to estimate the Wasserstein distance under the spiked transport model requires an exponential number of samples. By contrast, Section 6 establishes that, information theoretically, only a polynomial number of samples are required.

We now state our main result.

THEOREM 12. *There exists a positive universal constant c such that, for any d , estimating $W_1(\mu^{(1)}, \mu^{(2)})$ for distributions $\mu^{(1)}$ and $\mu^{(2)}$ on \mathbb{R}^d satisfying the spiked transport assumption with $k = 1$ to accuracy $\Theta(1/\sqrt{d})$ with probability at least $2/3$ requires at least 2^{cd} queries to $\text{VSTAT}(2^{cd})$.*

Our proof is based on a construction due to [26] [see also 14]. We defer the details to the appendix.

APPENDIX A: PROOFS OF LOWER BOUNDS

A.1 Proof of Theorem 4

We reduce from the spiked covariance model. By homogeneity, we may assume that $\sigma = 1$. Let $\mu^{(1)}$ be the standard Gaussian measure on \mathbb{R}^d , and for $\mu^{(2)}$ we take either the standard Gaussian measure or the distribution of a centered Gaussian with covariance $I + \beta vv^\top$, where $\|v\| = 1$ and $\beta > 0$ is to be specified. As long as $\beta \lesssim 1$, the measure $\mu^{(2)}$ is a $O(1)$ -Lipschitz pushforward of the Gaussian measure. Hence, it satisfies $T_p(O(1))$ [37, Corollary 20].

Note that if $\mu^{(2)}$ has covariance $I + \beta vv^\top$, then

$$W_p(\mu^{(1)}, \mu^{(2)}) \geq W_1(\mu^{(1)}, \mu^{(2)}) = W_1(\mathcal{N}(0, 1), \mathcal{N}(0, 1 + \beta)) \gtrsim \beta.$$

However, Cai et al. [15, Proposition 2] establish that the minimax testing error for

$$H_0 : \mathcal{N}(0, I) \text{ vs. } H_1 : \mathcal{N}(0, I + \beta vv^\top), \|v\| = 1$$

is bounded below by a constant when $\beta \lesssim \sqrt{d/n}$. A standard application of Le Cam's two-point method [84] yields the claim.

A.2 Proof of Proposition 10

We require the existence of two distributions on \mathbb{R}_+ , which will serve as the building blocks of our construction.

PROPOSITION 11. *For any integer $L \geq 0$ and $\varepsilon \in [0, 1/6]$, there exists a pair of random variables U and V with the following properties:*

- $\mathbb{E}U^j = \mathbb{E}V^j \quad \forall j \leq L$
- $U, V \in [0, 16\varepsilon^{-1}L^2]$ almost surely
- $\mathbb{E}U = \mathbb{E}V = 1$ and $\mathbb{E}U^2 = \mathbb{E}V^2 \leq 6$.
- $\mathbb{E}|U - 1| \leq 12\varepsilon$ but $\mathbb{E}|V - 1| \geq 1$.

The proof is deferred to Appendix A.3.

PROOF OF PROPOSITION 10. The proof follows closely the approach of Wu and Yang [92, Proposition 1]. We first employ a standard argument showing that we can consider the *Poissonalized* setting. It is trivial to see that given samples X_1, \dots, X_n from a distribution q on $[m]$, the counts $N_i = N_i(X_1, \dots, X_n) := |\{j \in [n] : X_j = i\}|$ are sufficient for q . We therefore consider tests ψ based on count vectors. Note that, under q , the count vector (N_1, \dots, N_m) has distribution Multinomial(n, q).

Define

$$R_n := \inf_{\psi} \left\{ \sup_{q \in \mathcal{D}_m^+} \mathbb{P}_q[\psi = 1] + \sup_{q \in \mathcal{D}_{m,\delta}^-} \mathbb{P}_q[\psi = 0] \right\}.$$

We aim to prove a lower bound on R_n .

Let $\rho > 0$, and let for $n \geq 1$ let $\{\psi_n\}$ be a set of near optimal tests for a fixed sample size; i.e.

$$\sup_{q \in \mathcal{D}_{m,\delta}^- \cup \mathcal{D}_m^+} \mathbb{P}_q[\psi_n \neq \mathbf{1}\{q \in \mathcal{D}_{m,\delta}^-\}] \leq R_n + \rho.$$

Define set of approximate probability vectors

$$\begin{aligned} \tilde{\mathcal{D}}_{m,\delta}^- &:= \left\{ q \in \mathbb{R}_+^m : \left| \sum_{i=1}^m q_i - 1 \right| \leq \delta, \frac{q}{\sum_{i=1}^m q_i} \in \mathcal{D}_{m,\delta}^- \right\} \\ \tilde{\mathcal{D}}_m^+ &:= \left\{ q \in \mathbb{R}_+^m : \left| \sum_{i=1}^m q_i - 1 \right| \leq \delta, \frac{q}{\sum_{i=1}^m q_i} \in \mathcal{D}_m^+ \right\}. \end{aligned}$$

We let $\tilde{\mathcal{D}}_{m,\delta} := \tilde{\mathcal{D}}_{m,\delta}^- \cup \tilde{\mathcal{D}}_m^+$. Given $q \in \tilde{\mathcal{D}}_{m,\delta}$, define the renormalization $\bar{q} = \sum_{i=1}^m q_i / q$. We then define

$$\tilde{R}_n := \inf_{\psi} \left\{ \sup_{q \in \tilde{\mathcal{D}}_m^+} \mathbb{P}_q[\psi = 1] + \sup_{q \in \tilde{\mathcal{D}}_{m,\delta}^-} \mathbb{P}_q[\psi = 0] \right\},$$

where the infimum is taken over all estimators based on the counts N_1, \dots, N_m and where \mathbb{P}_q indicates the probability when N_1, \dots, N_m are independent and $N_i \sim \text{Pois}(nq_i)$ for all $i \in [m]$. We set $N = \sum_{i=1}^m N_i$, and note that, conditioned on $N = n'$, the count vector (N_1, \dots, N_m) has distribution Multinomial(n', \bar{q}).

We define a test $\tilde{\psi}$ based on these Poissonalized counts by setting

$$\tilde{\psi}(N_1, \dots, N_m) := \psi_N(N_1, \dots, N_m).$$

This definition along with the near optimality of $\psi_{n'}$ for $n' \geq 0$ implies

$$\begin{aligned} \sup_{q \in \tilde{\mathcal{D}}_m^+} \mathbb{P}_q[\tilde{\psi} = 1] + \sup_{q \in \tilde{\mathcal{D}}_{m,\delta}^-} \mathbb{P}_q[\tilde{\psi} = 0] &\leq \sum_{n' \geq 0} R_{n'} \mathbb{P}_q[N = n'] + \rho \\ &\leq R_{n/2} + \mathbb{P}_q[N < n/2] + \rho, \end{aligned}$$

where the last inequality follows from the fact that $R_{n'} \leq 1$ for all $n' \geq 0$ and $R_{n'}$ is non-increasing in n' . Since $N = \text{Pois}(n \sum_{i=1}^m q_i)$ and $\sum_{i=1}^m q_i \geq 3/4$, a standard Chernoff bound implies $\mathbb{P}[N < n/2] \leq \exp(-Cn)$. Since ρ was arbitrary, we obtain that

$$\tilde{R}_n \leq R_{n/2} + \exp(-Cn).$$

To prove a lower bound on \tilde{R}_n , we consider random vectors

$$\begin{aligned} Q &= \frac{1}{m}(U_1, \dots, U_m) \\ Q' &= \frac{1}{m}(V_1, \dots, V_m), \end{aligned}$$

where U_i and V_i for $i \in [m]$ are independent copies of U and V constructed in Proposition 11 with $\varepsilon = \frac{1}{24}\delta$. Conditioned on Q and Q' , let N and N' be count vectors with independent entries generated by $N_i \sim \text{Pois}(nQ_i)$ and $N'_i \sim \text{Pois}(nQ'_i)$. Let us denote by \mathbb{P} and \mathbb{P}' the distributions of N and N' respectively. Under \mathbb{P} and \mathbb{P}' , the entries of N and N' are i.i.d. Poisson mixtures, so applying Wu and Yang [92, Lemma 4] yields

$$d_{\text{TV}}(\mathbb{P}, \mathbb{P}') \leq m \left(\frac{8enL}{\varepsilon m} \right)^L.$$

Let $E = \{Q \in \tilde{\mathcal{D}}_{m,\delta}^-\}$ and $E' = \{Q' \in \tilde{\mathcal{D}}_m^+\}$. By Lemma 10, $\mathbb{P}[E^C]$ and $\mathbb{P}[E'^C]$ are each at most $C \frac{L^4}{\delta^2 m}$.

Let π_E be the law of Q conditioned on E , and define $\pi'_{E'}$ analogously, and let \mathbb{P}_E and $\mathbb{P}'_{E'}$ be the laws of N and N' under these priors. We obtain for any estimator ψ based on count vectors

$$\begin{aligned} \sup_{q \in \tilde{\mathcal{D}}_m^+} \mathbb{P}_q[\tilde{\psi} = 1] + \sup_{q \in \tilde{\mathcal{D}}_{m,\delta}^-} \mathbb{P}_q[\tilde{\psi} = 0] &\geq \int \mathbb{P}_{q'}[\psi = 1] d\pi'_{E'}(q') + \int \mathbb{P}_q[\psi = 0] d\pi_E(q) \\ &\geq 1 - d_{\text{TV}}(\mathbb{P}_E, \mathbb{P}'_{E'}) \\ &\geq 1 - d_{\text{TV}}(\mathbb{P}, \mathbb{P}') - C \frac{L^4}{\delta^2 m}. \end{aligned}$$

Choosing $L = c \frac{\delta m}{n}$ for a sufficiently small constant c yields that

$$\tilde{R}_n \geq 1 - m \exp\left(-C \frac{\delta m}{n}\right) - C \frac{\delta^2 m^3}{n^4}.$$

Therefore

$$R_n \geq 1 - m \exp(-C \frac{\delta m}{n}) - C \frac{\delta^2 m^3}{n^4} - \exp(-Cn),$$

and choosing $m = \lceil C\delta^{-1}n \log n \rceil$ for C a sufficiently large constant and n sufficiently large yields the claim. \square

A.3 Proof of Proposition 11

First, the reduction of Wu and Yang [92, Lemma 7] implies that it suffices to construct random variables Y and Y' such that

- $\mathbb{E}Y^j = \mathbb{E}Y'^j \quad \forall 0 \leq j < L$
- $Y, Y' \in [1, 16\varepsilon^{-1}L^2]$ a.s.
- $\mathbb{E}Y = \mathbb{E}Y' \leq 6$
- $\mathbb{E}\frac{1}{Y} \geq 1 - 6\varepsilon$ but $\mathbb{E}\frac{1}{Y'} \leq \frac{1}{2}$.

Indeed, applying their construction yields U and V satisfying the first three requirements of Proposition 11 as well as $\mathbb{P}[U = 0] \leq 6\varepsilon$ and $\mathbb{P}[V = 0] \geq \frac{1}{4}$. Since the supports of U and V lie in $\{0\} \cup [1, +\infty)$, we have $\mathbb{E}|U - 1| = 2\mathbb{P}[U = 0] \leq 12\varepsilon$ and $\mathbb{E}|V - 1| = 2\mathbb{P}[V = 0] \geq 1$, as desired. We therefore focus on constructing such a Y and Y' .

By Wu and Yang [92, Lemma 7] combined with Timan [83, Section 2.11.1], there exist random variables X and X' supported on $[1, 16L^2]$ such that $\mathbb{E}X^j = \mathbb{E}X'^j$ for $0 \leq j < L$ and $\mathbb{E}\frac{1}{X} - \mathbb{E}\frac{1}{X'} \geq \frac{1}{2}$. Let P_ε and P'_ε denote the distribution of $\varepsilon^{-1}X$ and $\varepsilon^{-1}X'$, respectively.

Let

$$\begin{aligned} \Delta_\varepsilon &:= \int \frac{1}{(y-1)(y-2)} dP_\varepsilon(y) \\ \Delta'_\varepsilon &:= \int \frac{1}{(y'-1)(y'-2)} dP'_\varepsilon(y') \\ \mathcal{Z}_\varepsilon &:= \int \frac{1}{y-2} dP_\varepsilon(y) - \int \frac{1}{y'-1} dP'_\varepsilon(y'). \end{aligned}$$

We define two new distributions Q and Q' by

$$(6) \quad Q(dy) = \delta_1(dy) + \frac{1}{\mathcal{Z}_\varepsilon} \left(\frac{1}{(y-1)(y-2)} P_\varepsilon(dy) - \Delta_\varepsilon \delta_1(dy) \right)$$

$$(7) \quad Q'(dy') = \delta_2(dy') + \frac{1}{\mathcal{Z}_\varepsilon} \left(\frac{1}{(y'-1)(y'-2)} P'_\varepsilon(dy') - \Delta'_\varepsilon \delta_2(dy') \right)$$

By Lemma 8,

$$\begin{aligned} \Delta_\varepsilon, \Delta'_\varepsilon &\in [0, \frac{9}{5}\varepsilon^2] \\ \mathcal{Z}_\varepsilon &\geq \frac{3}{10}\varepsilon, \end{aligned}$$

which implies in particular that both Q and Q' are probability distributions.

Let $Y \sim Q$ and $Y' \sim Q'$. We first check the last three conditions. Clearly Y and Y' are supported on $[1, 16\varepsilon^{-1}L^2]$, and Lemma 9 implies that $\mathbb{E}Y = \mathbb{E}Y' \leq 6$. We have $\mathbb{E}\frac{1}{Y} \geq \mathbb{P}[Y = 1] = 1 - \frac{\Delta_\varepsilon}{\mathcal{Z}_\varepsilon} \geq 1 - 6\varepsilon$, and since $Y' \geq 2$ almost surely the bound $\mathbb{E}\frac{1}{Y'} \leq \frac{1}{2}$ is immediate.

It remains to check the moment-matching condition. Any polynomial $p(y)$ of degree at most $L - 1$ can be written

$$p(y) = (y - 1)(y - 2)q(y) + \alpha y + \beta,$$

where $q(y)$ has degree less than $L - 1$. Then

$$\mathbb{E}p(Y) - \mathbb{E}p(Y') = \mathbb{E}(Y - 1)(Y - 2)q(Y) - \mathbb{E}(Y' - 1)(Y' - 2)q(Y') + \alpha(\mathbb{E}Y - \mathbb{E}Y').$$

The last term vanishes because $\mathbb{E}Y = \mathbb{E}Y'$, and

$$\mathbb{E}(Y - 1)(Y - 2)q(Y) - \mathbb{E}(Y' - 1)(Y' - 2)q(Y') = \frac{1}{Z_\varepsilon} (\mathbb{E}q(\varepsilon^{-1}X) - \mathbb{E}q(\varepsilon^{-1}X')) = 0,$$

since $\mathbb{E}X^j = \mathbb{E}X'^j$ for all $j < L - 1$.

Therefore $\mathbb{E}p(Y) = \mathbb{E}p(Y')$ for all polynomials of degree at most $L - 1$. \square

A.4 Proof of Theorem 12

We first establish the existence of a probability distribution on \mathbb{R} which agrees with $\mathcal{N}(0, 1)$ on many moments, but is far from $\mathcal{N}(0, 1)$ in Wasserstein distance.

PROPOSITION 12. *There exists a $O(1)$ -subgaussian distribution A on \mathbb{R} that satisfies the following requirements.*

- A agrees with $\mathcal{N}(0, 1)$ on the first $2m - 1$ moments.
- $W_1(A, \mathcal{N}(0, 1)) = \Omega(1/\sqrt{m})$.
- $\chi^2(A, \mathcal{N}(0, 1)) = \exp(O(m))$.

PROOF. By Diakonikolas et al. [26, Corollary 4.4], for any $\delta \in (0, 1)$ we can find an atomic measure Q supported on m points in $[-O(\sqrt{m}), O(\sqrt{m})]$ such that $A := Q * \mathcal{N}(0, \delta)$ matches $\mathcal{N}(0, 1)$ on the first $2m - 1$ moments. By Diakonikolas et al. [26, Lemma 4.6], this distribution satisfies the bound $\chi^2(A, \mathcal{N}(0, 1)) = \exp(O(m))/\sqrt{\delta}$. Moreover, the distribution Q is supported on the zeros of a rescaled version of the m th Hermite polynomial, and by Bubeck et al. [14, Lemma A.1] these zeros are $\Omega(1/\sqrt{m})$ apart.

Since the points in the support of Q are $\Omega(1/\sqrt{m})$ apart, there exists a constant c such that the union of balls of radius c/\sqrt{m} centered at the support of Q covers at most half of the interval $[-1, 1]$, and since the Gaussian density is bounded below on this interval, a constant fraction of the mass of $\mathcal{N}(0, 1)$ is located at distance at least c/\sqrt{m} from Q . Hence $W_1(Q, \mathcal{N}(0, 1)) = \Omega(1/\sqrt{m})$. Clearly $W_1(A, Q) = O(\sqrt{\delta})$.

Therefore, if we choose $\delta = O(1/m)$, then

$$W_1(A, \mathcal{N}(0, 1)) \geq W_1(Q, \mathcal{N}(0, 1)) - W_1(A, Q) = \Omega(1/\sqrt{m})$$

and $\chi^2(A, \mathcal{N}(0, 1)) = O(\sqrt{m}) \exp(O(m)) = \exp(O(m))$.

Finally, we show that A is $O(1)$ -subgaussian, and therefore satisfies $T_1(C)$ for a positive constant C . Standard facts [see 88] imply that it suffices to show that if $X \sim A$, then $\|X\|_k = O(\sqrt{k})$ for all k . This clearly holds for $k \leq 2m - 1$, since $\mathcal{N}(0, 1)$ is itself 1-subgaussian and the first $2m - 1$ moments of A and

$\mathcal{N}(0,1)$ agree. On the other hand, for $k \geq 2m$, if $Y \sim Q$ and $Z \sim \mathcal{N}(0,1)$, then $X = Y + \sqrt{\delta}Z \sim A$, and

$$\|X + \sqrt{\delta}Z\|_k \leq \|X\|_\infty + \sqrt{\delta}\|Z\|_k \lesssim \sqrt{m} + \sqrt{k/m} \lesssim \sqrt{k},$$

as desired. \square

The separation $W_1(A, \mathcal{N}(0,1)) = \Omega(1/\sqrt{m})$ in Proposition 12 is easily seen to be tight. Indeed, Rigollet and Weed [72, Corollary 2] show that if μ and ν are $O(1)$ -subgaussian and agree on their first $O(m)$ moments, then $W_1(\mu, \nu) = O(1/\sqrt{m})$.

By planting the distribution constructed in Proposition 12 in a random direction, we obtain two high-dimensional measures satisfying the spiked transport model.

LEMMA 3. *Let v be a unit vector in \mathbb{R}^d , and denote by P_v the distribution on \mathbb{R}^d of the random variable $Xv + Z$, where $X \sim A$ and $Z \sim \mathcal{N}(0, I_d - vv^\top)$ is independent of X . Then $\mu^{(1)} = P_v$ and $\mu^{(2)} = \mathcal{N}(0, I_d)$ satisfy the spiked transport model (2), and $W_1(P_v, \mathcal{N}(0, I_d)) = \Omega(1/\sqrt{m})$.*

PROOF. If we let $\xi \sim \mathcal{N}(0,1)$, then $\mathcal{N}(0, I_d)$ is the law of $\xi v + Z$, where ξ and Z are independent and $Z \sim \mathcal{N}(0, I_d - vv^\top)$. Denoting by \mathcal{U} the span of v , we see that P_v and $\mathcal{N}(0, I_d)$ satisfy (2) with $X^{(1)} = Xv$ and $X^{(2)} = \xi v$. By Propositions 3 and 12, $W_1(P_v, \mathcal{N}(0, I_d)) = W_1(A, \mathcal{N}(0,1)) = \Omega(1/\sqrt{m})$. \square

The proof of Theorem 12 follows from a framework due to Feldman et al. [30], from which the following result is extracted. Given distributions P_1, P_2 , and Q , define

$$\chi_Q^2(P_1, P_2) := \int \left(\frac{dP_1}{dQ} - 1 \right) \left(\frac{dP_2}{dQ} - 1 \right) dQ.$$

We call a set \mathcal{P} of distributions (γ, β) correlated with respect to Q if for all $P_i, P_j \in \mathcal{P}$,

$$\chi_Q^2(P_i, P_j) \leq \begin{cases} \beta & \text{if } i = j, \\ \gamma & \text{if } i \neq j. \end{cases}$$

We then have the following.

PROPOSITION 13. *Let \mathcal{Q} be a set of distributions, and let Q be a reference distribution. Suppose that there exists a set $\mathcal{P} \subseteq \mathcal{Q}$ such that \mathcal{P} is (γ, β) correlated with respect to Q . Then any SQ algorithm that distinguishes queries from $P = Q$ and $P \in \mathcal{P}$ requires at least $|\mathcal{P}|\gamma/3\beta$ queries to VSTAT($1/2\gamma$)*

PROOF. By choosing $\gamma' = \gamma$ in Feldman et al. [30, Lemma 3.10], we obtain that the set \mathcal{P} satisfies $\text{SDA}(\mathcal{P}, Q, 2\gamma) \geq |\mathcal{P}|\gamma/(\beta - \gamma) \geq |\mathcal{P}|\gamma/\beta$. Then Feldman et al. [30, Theorem 3.7] implies that distinguishing Q from \mathcal{P} requires at least $|\mathcal{P}|\gamma/3\beta$ queries to VSTAT($1/2\gamma$). \square

We can now prove the lower bound.

PROOF OF THEOREM 12. Let \mathcal{Q} be the set $\{P_v : v \in \mathbb{R}^d, \|v\| = 1\}$. The Johnson-Lindenstrauss lemma [46] implies that for any $\delta \in (0,1)$, there exists a

set of $2^{\Omega(\delta^2 d)}$ unit vectors in \mathbb{R}^d with pairwise inner product at most δ . Denote by \mathcal{S} a set of such vectors, and set $\mathcal{P} := \{P_v : v \in \mathcal{S}\} \subseteq \mathcal{Q}$.

Write Δ for $\chi^2(A, \mathcal{N}(0, 1))$, and recall that $\Delta = \exp(O(m))$. By Diaconikolas et al. [26, Lemma 3.4], $\chi_{\mathcal{N}(0, I)}^2(P_v, P_{v'}) \leq |v \cdot v'|^{2m} \Delta$, and the set \mathcal{P} is therefore $(\delta^{2m} \Delta, \Delta)$ correlated. By Proposition 13, any SQ algorithm that distinguishes between $\mu = \mathcal{N}(0, 1)$ and $\mu \in \mathcal{P}$ with probability at least $2/3$ requires $2^{\Omega(\delta^2 d)} \delta^{2m}$ queries to $\text{VSTAT}(\delta^{-2m}/2\Delta)$.

Let $\delta > 0$ be a constant small enough that $\delta^{-2m}/2\Delta = \exp(\Omega(m))$. If $m = cd$ for a sufficiently small positive constant c , then $2^{\Omega(\delta^2 d)} \delta^{2m} = \exp(\Omega(d))$. An SQ algorithm to distinguish queries from $\mathcal{N}(0, 1)$ from those from a distribution in \mathcal{Q} with probability at least $2/3$ therefore requires $\exp(\Omega(d))$ queries to $\text{VSTAT}(\exp(\Omega(d)))$. Therefore, by Lemma 3, any SQ algorithm which estimates W_1 under the spiked transport model to accuracy $\Theta(1/\sqrt{m}) = \Theta(1/\sqrt{d})$ requires $\exp(\Omega(d))$ queries to $\text{VSTAT}(\exp(\Omega(d)))$, as claimed. \square

APPENDIX B: SUPPLEMENTARY MATERIAL

B.1 Additional lemmas

LEMMA 4. *There exists a universal constant c such that for all $\varepsilon \in (0, 1)$, the covering number of $\mathcal{V}_k(\mathbb{R}^d)$ satisfies*

$$\log \mathcal{N}(\mathcal{V}_k, \varepsilon, \|\cdot\|_{\text{op}}) \leq dk \log \frac{c\sqrt{k}}{\varepsilon}.$$

PROOF. If we denote by \mathcal{S}_d the unit sphere in \mathbb{R}^d , then $\mathcal{V}_k(\mathbb{R}^d) \subseteq \mathcal{S}_d^{\times k}$, and

$$\|U - V\|_{\text{op}}^2 \leq \sum_{i=1}^k \|U_i - V_i\|^2,$$

where $\{U_i\}$ (resp. $\{V_i\}$) are the rows of U (resp. V). The claim follows immediately from the fact that $\log \mathcal{N}(\mathcal{S}_d, \varepsilon, \|\cdot\|) \leq d \log \frac{c}{\varepsilon}$ for a universal constant c . \square

LEMMA 5. *Let μ be a centered distribution satisfying $T_1(\sigma^2)$. If $X \sim \mu$, then $(\mathbb{E}\|X\|^p)^{1/p} \lesssim \sqrt{dp}$.*

PROOF. By the monotonicity of L_p norms, we may assume that p is a positive even integer. The claim then follows from Proposition 16 and standard bounds on the moments of subgaussian random variables [88]. \square

The following proposition provides a sort of converse to Proposition 3: if the k -dimensional Wasserstein distance agrees with the Wasserstein distance, then there is an optimal coupling between the measures which acts only on a k -dimensional subspace.

PROPOSITION 14. *If $\tilde{W}_{p,k}(\mu, \nu) = W_p(\mu, \nu)$, then there exists a coupling γ optimal for μ and ν such that $X - Y$ is supported on a k -dimensional subspace.*

PROOF. Let γ be an optimal coupling for μ and ν , and write

$$U = \operatorname{argmax}_{V \in \mathcal{V}_k(\mathbb{R}^d)} W_p(\mu_V, \nu_V).$$

Then

$$W_p^p(\mu, \nu) = \int \|x - y\|^p d\gamma(x, y) \geq \int \|U(x - y)\|^p d\gamma(x, y) \geq \tilde{W}_{p,k}^p(\mu, \nu),$$

where the first inequality uses the fact that $\|U\|_{\text{op}} = 1$ and the second uses the fact that if $(X, Y) \sim \gamma$, then UX and UY are distributed according to μ_U and ν_U , respectively. If the first equality holds, then we must have $\|X - Y\| = \|U(X - Y)\|$ γ -almost surely, and if the second inequality holds then this coupling is optimal for μ_U and ν_U \square

PROPOSITION 15 (88, Proposition 2.5.2). *If a real-valued random variable X is σ^2 -subgaussian, then*

$$(8) \quad \mathbb{E}e^{\lambda^2 X^2} \leq e^{4\lambda^2 \sigma^2} \quad \forall |\lambda| \leq \frac{1}{2\sigma}.$$

Conversely, if (8) holds and $\mathbb{E}X = 0$, then X is $8\sigma^2$ -subgaussian.

PROOF. For the first claim, let Z be a standard Gaussian random variable independent of X . Then, if $\lambda^2 \sigma^2 < 1/4$,

$$\mathbb{E}e^{\lambda^2 X^2} = \mathbb{E}e^{\sqrt{2}\lambda Z X} \leq \mathbb{E}e^{\lambda^2 \sigma^2 Z^2} = \frac{1}{(1 - 2\lambda^2 \sigma^2)^{1/2}} \leq e^{4\lambda^2 \sigma^2},$$

where the last step uses the inequality $(1 - x)^{-1} \leq e^{2x}$ for $0 \leq x \leq \frac{1}{2}$.

Conversely, suppose that (8) holds and $\mathbb{E}X = 0$. Then, for $|\lambda| \leq \frac{1}{2\sigma}$, we have

$$\mathbb{E}e^{\lambda X} \leq \mathbb{E}\lambda X + \mathbb{E}e^{\lambda^2 X^2} \leq e^{4\lambda^2 \sigma^2}.$$

If $|\lambda| > \frac{1}{2\sigma}$, then

$$\mathbb{E}e^{\lambda X} \leq \mathbb{E}e^{2\lambda^2 \sigma^2 + \frac{X^2}{8\sigma^2}} \leq e^{2\lambda^2 \sigma^2 + 1} \leq e^{4\lambda^2 \sigma^2}.$$

\square

PROPOSITION 16. *If X on \mathbb{R}^k is σ^2 -subgaussian, then $\varepsilon\|X\|$ is $(8k\sigma^2)$ -subgaussian, where ε is a Rademacher random variable independent of X .*

PROOF. For $|\lambda| < \frac{1}{2\sqrt{k}\sigma}$,

$$\mathbb{E}\exp(\lambda^2(\varepsilon\|X\|)^2) = \mathbb{E}\exp(\lambda^2\|X\|^2) \leq \mathbb{E}\exp(k\lambda^2(e_i^\top X)^2) \leq \exp(4k\lambda^2\sigma^2).$$

Applying Proposition 15 yields the claim. \square

PROPOSITION 17. *Let $p' \in [1, 2]$. If ρ_n on \mathbb{R}^k satisfies $T_{p'}(\sigma^2)$, then for any $p \in [1, \infty)$,*

$$\mathbb{E}W_p(\rho_n, \rho) \leq r_{p,k}(n) := c_p \sigma \sqrt{k} \begin{cases} n^{-1/2p} & \text{if } k < 2p \\ n^{-1/2p}(\log n)^{1/p} & \text{if } k = 2p \\ n^{-1/k} & \text{if } k > 2p \end{cases}$$

PROOF. Lemma 1 implies that $X \sim \rho$ is σ^2 subgaussian. in particular, X satisfies $(\mathbb{E}\|X\|^q)^{1/q} \leq \sigma\sqrt{qk}$. Choosing $q = 3p$ and using Lei [57, Theorem 3.1] implies the claim. \square

LEMMA 6. *If $\mathcal{N}_\varepsilon(\mathcal{X}) \geq c\varepsilon^{-d}$, then for any positive integer m there exists a subset of \mathcal{X} of cardinality m such that each pair of points is separated by at least $\frac{1}{2}\left(\frac{c}{m}\right)^{1/d}$.*

PROOF. Let S be a maximal subset of \mathcal{X} such that each pair of points in S is separated by at least $\frac{1}{2}\left(\frac{c}{m}\right)^{1/d}$. Then, for any point $x \in X$, there must be an $s \in S$ such that $d(x, s) < \frac{1}{2}\left(\frac{c}{m}\right)^{1/d}$, since otherwise x could be added to S . Therefore, balls of radius $\frac{1}{2}\left(\frac{c}{m}\right)^{1/d}$ around the points $s \in S$ cover X , which implies

$$|S| \geq \mathcal{N}\left(X, \left(\frac{c}{m}\right)^{1/d}\right) \geq m,$$

as claimed. \square

LEMMA 7. *Let Q_1, \dots, Q_ℓ be a partition of \mathcal{X} . If F is a uniform random bijection from $[m]$ to \mathcal{G}_m , then*

$$\mathbb{E} \sum_{i=1}^{\ell} |F_{\#}q(Q_i) - F_{\#}u(Q_i)| \leq \|q - u\|_1 \wedge \sqrt{\frac{\ell \cdot \chi^2(q, u)}{m}}$$

PROOF. The first bound follows immediately from the triangle inequality:

$$\sum_{i=1}^{\ell} |F_{\#}q(Q_i) - F_{\#}u(Q_i)| \leq \sum_{i=1}^{\ell} \sum_{j:F(j) \in Q_i} \left| q(j) - \frac{1}{m} \right| = \|q - u\|_1.$$

We now show the second bound. Fix $i \in [\ell]$. Under the random choice of F , the quantity $F_{\#}q(Q_i) - F_{\#}u(Q_i) = \sum_{j:F(j) \in Q_i} \left(q(j) - \frac{1}{m}\right)$ is a sum of $|Q_i|$ terms selected uniformly *without replacement* from the set $\Delta := \{q(j) - \frac{1}{m} : j \in [m]\}$. By Hoeffding [42, Theorem 4], since $x \mapsto |x|$ is continuous and convex,

$$\mathbb{E} \left| \sum_{j:F(j) \in Q_i} \left(q(j) - \frac{1}{m}\right) \right| \leq \mathbb{E} \left| \sum_{j=1}^{|Q_i|} \Delta_j \right|,$$

where Δ_j are selected uniformly *with replacement* from Δ . Note that $\mathbb{E}\Delta_j = 0$, and $\mathbb{E}\Delta_j^2 = \frac{1}{m} \sum_{j=1}^m \left(q(j) - \frac{1}{m}\right)^2 = m^{-2} \chi^2(q, u)$. Jensen's inequality then yields

$$\mathbb{E} \left| \sum_{j=1}^{|Q_i|} \Delta_j \right| \leq \frac{\sqrt{|Q_i| \chi^2(q, u)}}{m}.$$

The Cauchy-Schwarz inequality finally implies

$$\mathbb{E} \sum_{i=1}^{\ell} \left| \sum_{j:F(j) \in Q_i} \left(q(j) - \frac{1}{m}\right) \right| \leq \sum_{i=1}^{\ell} \frac{\sqrt{|Q_i| \chi^2(q, u)}}{m} \leq \sqrt{\frac{\ell \cdot \chi^2(q, u)}{m}}.$$

\square

LEMMA 8. Let $\varepsilon \in [0, 1/6]$. If $X, X' \geq 1$ almost surely and $\mathbb{E}\frac{1}{X} - \mathbb{E}\frac{1}{X'} \geq \frac{1}{2}$, then

$$\mathbb{E}\frac{1}{(\varepsilon^{-1}X - 1)(\varepsilon^{-1}X - 2)}, \mathbb{E}\frac{1}{(\varepsilon^{-1}X' - 1)(\varepsilon^{-1}X' - 2)} \in [0, \frac{9}{5}\varepsilon^2]$$

and

$$\mathbb{E}\frac{1}{\varepsilon^{-1}X - 2} - \mathbb{E}\frac{1}{\varepsilon^{-1}X' - 1} \geq \frac{3}{10}\varepsilon.$$

PROOF. If $x \geq 1$ and $\varepsilon \leq 1/6$, then $(x - \varepsilon)(x - 2\varepsilon) \geq \frac{5}{9}$, so

$$\frac{1}{(\varepsilon^{-1}x - 1)(\varepsilon^{-1}x - 2)} = \frac{\varepsilon^2}{(x - \varepsilon)(x - 2\varepsilon)} \leq \frac{9}{5}\varepsilon^2,$$

and since $X, X' \geq 1$ almost surely the first claim is immediate.

Similarly, if $x' \geq 1$ and $\varepsilon \leq 1/6$, we have $\frac{\varepsilon}{x'(x' - \varepsilon)} \leq \frac{1}{5}$; hence almost surely

$$\frac{1}{\varepsilon^{-1}X - 2} - \frac{1}{\varepsilon^{-1}X' - 1} \geq \frac{\varepsilon}{X} - \frac{\varepsilon}{X'} - \frac{\varepsilon^2}{X'(X' - \varepsilon)} \geq \varepsilon \left(\frac{1}{X} - \frac{1}{X'} - \frac{1}{5} \right),$$

and taking expectations yields the claim. \square

LEMMA 9. Let Q and Q' be defined as in (6) and (7). If $Y \sim Q$ and $Y' \sim Q'$, then

$$\mathbb{E}Y = \mathbb{E}Y' \leq 6.$$

PROOF. It follows directly from the definition that

$$\begin{aligned} \mathbb{E}Y &= 1 + \frac{1}{\mathcal{Z}_\varepsilon} \left(\int \frac{y}{(y-1)(y-2)} P_\varepsilon(dy) - \Delta_\varepsilon \right) \\ &= 1 + \frac{1}{\mathcal{Z}_\varepsilon} \int \frac{1}{y-2} P_\varepsilon(dy) \end{aligned}$$

and analogously

$$\begin{aligned} \mathbb{E}Y' &= 2 + \frac{1}{\mathcal{Z}_{\varepsilon'}} \left(\int \frac{y'}{(y'-1)(y'-2)} P_{\varepsilon'}(dy') - 2\Delta_{\varepsilon'} \right) \\ &= 2 + \frac{1}{\mathcal{Z}_{\varepsilon'}} \int \frac{1}{y'-1} P_{\varepsilon'}(dy'), \end{aligned}$$

which implies

$$\begin{aligned} \mathbb{E}Y' - \mathbb{E}Y &= 1 - \frac{1}{\mathcal{Z}_\varepsilon} \left(\int \frac{1}{y-2} dP_\varepsilon(y) - \int \frac{1}{y'-1} dP'_{\varepsilon'}(y') \right) \\ &= 1 - \frac{1}{\mathcal{Z}_\varepsilon} (\mathcal{Z}_\varepsilon) = 0. \end{aligned}$$

To verify the inequality, note that

$$\int \frac{1}{y-2} P_\varepsilon(dy) = \mathbb{E}\frac{1}{\varepsilon^{-1}X - 2} = \varepsilon \mathbb{E}\frac{1}{X - 2\varepsilon},$$

and the fact that $X \geq 1$ almost surely and $\varepsilon \leq 1/6$ implies that this quantity is bounded by $\frac{3}{2}\varepsilon$. By Lemma 8, $\mathcal{Z}_\varepsilon \geq \frac{3}{10}\varepsilon$; therefore

$$\mathbb{E}Y = 1 + \frac{1}{\mathcal{Z}_\varepsilon} \int \frac{1}{y-2} P_\varepsilon(dy) \leq 1 + \frac{\frac{3}{2}\varepsilon}{\frac{3}{10}\varepsilon} = 6.$$

\square

LEMMA 10. *If $Q = \frac{1}{m}(U_1, \dots, U_m)$ and $Q' = \frac{1}{m}(V_1, \dots, V_m)$, where U and V satisfy the assumptions of Proposition 11 with $\varepsilon = \frac{1}{24}\delta$, and recall that $E = \{Q \in \tilde{\mathcal{D}}_{m,\delta}^-\}$ and $E' = \{Q \in \tilde{\mathcal{D}}_m^+\}$. Then*

$$\mathbb{P}[E^C] + \mathbb{P}[E'^C] \lesssim \frac{L^4}{\delta^2 m}.$$

PROOF. We first show that E holds as long as $\frac{1}{m} \sum_{i=1}^m U_i^2 \leq 8$ and $\frac{1}{m} \sum_{i=1}^m |U_i - 1| \leq \delta$. Indeed, if we set $s := \sum_{i=1}^m Q_i$, where $Q_i = U_i/m$, then the second condition together with Jensen's inequality imply that $|s - 1| \leq \delta$. Therefore, recalling that $\bar{Q} = Q/s$, we get

$$\chi^2(\bar{Q}, u) + 1 = m \sum_{i=1}^m \bar{Q}_i^2 = \frac{1}{s^2 m} \sum_{i=1}^m U_i^2 \leq \frac{8}{s^2} \leq \frac{8}{1 - 2\delta} \leq 10,$$

where the final inequality follows from the assumption that $\delta \leq 1/10$. Hence $\chi^2(\bar{Q}, u) \leq 9$.

Likewise,

$$d_{\text{TV}}(\bar{Q}, u) = \frac{1}{2} \sum_{i=1}^m \left| \bar{Q}_i - \frac{1}{m} \right| \leq \frac{1}{2} \sum_{i=1}^m |\bar{Q}_i - Q_i| + \frac{1}{2} \sum_{i=1}^m \left| Q_i - \frac{1}{m} \right| \leq \frac{1}{2} (|s-1| + \delta) \leq \delta.$$

By Chebyshev's inequality combined with the assumption that $\mathbb{E}U^2 \leq 6$,

$$\mathbb{P} \left[\frac{1}{m} \sum_{i=1}^m U_i^2 > 8 \right] \leq \frac{\text{Var}(U^2)}{4m},$$

and since $U \leq 16\varepsilon^{-1}L^2$ almost surely,

$$\text{Var}(U^2) \leq \mathbb{E}U^4 \leq (16\varepsilon^{-1}L^2)^2 \mathbb{E}U^2 \lesssim \delta^{-2}L^4.$$

Moreover, since $\varepsilon = \frac{1}{24}\delta$ and $\mathbb{E}|U - 1| \leq 12\varepsilon$,

$$\mathbb{P} \left[\frac{1}{m} \sum_{i=1}^m |U_i - 1| > \delta \right] \leq \frac{4 \text{Var}(|U - 1|)}{\delta^2 m} \lesssim \frac{1}{\delta^2 m}.$$

Combining these bounds yields

$$\mathbb{P}[E^C] \lesssim \frac{L^4}{\delta^2 m}.$$

The argument for E' is analogous. It suffices for E' to hold that $\frac{1}{m} \sum_{i=1}^m V_i^2 \leq 8$ along with the conditions $\frac{1}{m} \sum_{i=1}^m |V_i - 1| \geq \frac{3}{4}$ and $|(\frac{1}{m} \sum_{i=1}^m V_i) - 1| \leq \delta$. As above, the first condition is violated with probability at most $C \frac{L^4}{\delta^2 m}$, and the latter two conditions are violated with probability at most $C \frac{1}{\delta^2 m}$. The claim follows. \square

B.2 Omitted proofs

PROOF OF PROPOSITION 4. Denote by $\rho^{(1)}$ and $\rho^{(2)}$ the distributions of $Z^{(1)}$ and $Z^{(2)}$, respectively. It suffices to show $\tilde{W}_{p,k}(\nu^{(1)}, \nu^{(2)}) + W_p(\rho^{(1)}, \rho^{(2)}) \geq W_p(\nu^{(1)}, \nu^{(2)})$. Let $U \in \mathcal{V}_k(\mathbb{R}^d)$ be such that the column span of U agrees with the subspace \mathcal{U} on which the random variables $X^{(1)}$ and $X^{(2)}$ are supported in the definition of $\mu^{(1)}$ and $\mu^{(2)}$. We let $(Y^{(1)}, Y^{(2)})$ be a pair of random variables such that marginally $Y^{(i)} \sim \mu_U^{(i)}$ for $i \in \{1, 2\}$ and

$$W_p^p(\mu_U^{(1)}, \mu_U^{(2)}) = \mathbb{E}\|Y^{(1)} - Y^{(2)}\|^p.$$

Likewise, let $(W^{(1)}, W^{(2)})$ be a coupling of $\rho^{(1)}$ and $\rho^{(2)}$ such that

$$W_p^p(\rho^{(1)}, \rho^{(2)}) = \mathbb{E}\|W^{(1)} - W^{(2)}\|^p.$$

Then (4) implies that $(Y^{(1)} + W^{(1)}, Y^{(2)} + W^{(2)})$ is a coupling of $\mu^{(1)}$ and $\mu^{(2)}$. Therefore

$$\begin{aligned} W_p(\mu^{(1)}, \mu^{(2)}) &\leq \left(\mathbb{E}\|Y^{(1)} + W^{(1)} - (Y^{(2)} + W^{(2)})\|^p \right)^{1/p} \\ &\leq \left(\mathbb{E}\|Y^{(1)} - Y^{(2)}\|^p \right)^{1/p} + \left(\mathbb{E}\|W^{(1)} - W^{(2)}\|^p \right)^{1/p} \\ &= W_p(\mu_U^{(1)}, \mu_U^{(2)}) + W_p(\rho^{(1)}, \rho^{(2)}) \\ &\leq \tilde{W}_{p,k}(\mu^{(1)}, \mu^{(2)}) + W_p(\rho^{(1)}, \rho^{(2)}). \end{aligned}$$

□

PROOF OF LEMMA 1. This can be deduced from a series of well known facts. By Proposition 6, it suffices to show that for f Lipschitz,

$$\mathbb{E}e^{\lambda(f(X) - \mathbb{E}f(X))} \leq e^{C^2 k \lambda^2 \sigma^2 / 2}.$$

By symmetrization, we can bound

$$\mathbb{E}e^{\lambda(f(X) - \mathbb{E}f(X))} \leq \mathbb{E}e^{\lambda\varepsilon(f(X) - f(X'))} \leq \mathbb{E}e^{\lambda\varepsilon\|X - X'\|},$$

where $X' \sim \mu$ is independent of X . Continuing, we obtain

$$\mathbb{E}e^{\lambda\varepsilon\|X - X'\|} \leq \mathbb{E}e^{2\lambda\varepsilon\|X\|} \leq e^{16k\lambda^2\sigma^2}$$

by Proposition 16. This proves the claim with $C = 32$. □

PROOF OF THEOREM 6. First, assume that μ satisfies the $T_p(\sigma^2)$ inequality. Then, by Gozlan and Léonard [40, Proposition 1.9], the product measure $\mu^{\otimes n}$ satisfies the inequality $T_p(n^{2/p-1}\sigma^2)$ on the space \mathcal{X}^n equipped with the metric $\rho_p(x, y) := (\sum_{i=1}^n \rho(x_i, y_i)^p)^{1/p}$. Therefore, $\mu^{\otimes n}$ also satisfies $T_1(n^{2/p-1}\sigma^2)$ with respect to this same metric.

We now note that the function

$$(x_1, \dots, x_n) \mapsto W_p \left(\frac{1}{n} \sum_{i=1}^n \delta_{x_i}, \mu \right)$$

is $n^{-1/p}$ -Lipschitz with respect to ρ_p . Indeed, we have

$$\begin{aligned} \left| W_p\left(\frac{1}{n} \sum_{i=1}^n \delta_{x_i}, \mu\right) - W_p\left(\frac{1}{n} \sum_{i=1}^n \delta_{y_i}, \mu\right) \right| &\leq W_p\left(\frac{1}{n} \sum_{i=1}^n \delta_{x_i}, \frac{1}{n} \sum_{i=1}^n \delta_{y_i}\right) \\ &\leq \left(\frac{1}{n} \sum_{i=1}^n \rho(x_i, y_i)^p\right)^{1/p} \\ &= n^{-1/p} \rho_p(x, y). \end{aligned}$$

Combining these observations with Proposition 6, we obtain

$$\mathbb{E}e^{\lambda(W_p(\mu_n, \mu) - \mathbb{E}W_p(\mu_n, \mu))} \leq e^{\lambda^2 n^{-2/p} n^{2/p-1} \sigma^2 / 2} = e^{\lambda^2 n^{-1} \sigma^2 / 2} \quad \forall \lambda \in \mathbb{R}.$$

In the other direction, suppose that $W_p(\mu_n, \mu)$ is σ^2/n -subgaussian. A Chernoff bound implies the concentration inequality

$$\mathbb{P}[W_p(\mu_n, \mu) - \mathbb{E}W_p(\mu_n, \mu) \geq t] \leq e^{nt^2/2\sigma^2}.$$

We follow the proof of Gozlan [38, Theorem 3.4], and only sketch the argument here. It is easy to verify that the function $W_p(\cdot, \mu)$ is lower semi-continuous with respect to the weak topology [see 89, proof of Theorem 4.1], which implies that, for any $t \geq 0$, the set $\mathcal{O}_t := \{\nu \in \mathcal{P}(\mathcal{X}) : W_p(\nu, \mu) > t\}$ is open. Moreover, as $\mu \in \mathcal{P}_p$, Villani [89, Theorem 6.9] combined with Varadarajan's theorem yields that $\mathbb{E}W_p(\mu_n, \mu) \rightarrow 0$ as $n \rightarrow \infty$ [see 89, proof of Theorem 22.22].

Applying Sanov's theorem then yields

$$\begin{aligned} - \inf_{\nu \in \mathcal{O}_t} D(\nu \|\mu) &\leq \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}[W_p(\mu_n, \mu) > t] \\ &\leq \liminf_{n \rightarrow \infty} \frac{1}{n} \cdot \frac{n \max\{(t - \mathbb{E}W_p(\mu_n, \mu))^2, 0\}}{2\sigma^2} \\ &= \frac{t^2}{2\sigma^2}. \end{aligned}$$

We obtain that, for any $t \geq 0$,

$$D(\nu \|\mu) \geq -\frac{t^2}{2\sigma^2} \quad \forall \nu \in \mathcal{P}(\mathcal{X}) \text{ s.t. } W_p(\nu, \mu) > t.$$

Therefore μ satisfies $T_p(\sigma^2)$. □

PROOF OF THEOREM 7. As in the proof of Theorem 6, we have that $\mu^{\otimes n}$ satisfies $T_1(n^{2/p'-1}\sigma^2)$ with respect to $\rho_{p'}(x, y)$. We also that that

$$(x_1, \dots, x_n) \mapsto W_p\left(\frac{1}{n} \sum_{i=1}^n \delta_{x_i}, \mu\right)$$

is $n^{-1/p}$ -Lipschitz with respect to $\rho_{p'}$, since the fact that $p' \leq p$ implies

$$\left| W_p\left(\frac{1}{n} \sum_{i=1}^n \delta_{x_i}, \mu\right) - W_p\left(\frac{1}{n} \sum_{i=1}^n \delta_{y_i}, \mu\right) \right| \leq n^{-1/p} \rho_p(x, y) \leq n^{-1/p} \rho_{p'}(x, y).$$

We conclude as in the first part of the proof of Theorem 6. □

PROOF OF THEOREM 10. We show that for any measures $\nu^{(1)}, \nu^{(2)}$, if $\mathcal{V} = \text{span}(\hat{V})$ where $V = \text{argmax}_{U \in \mathcal{V}_k(\mathbb{R}^d)} W_p(\nu_U^{(1)}, \nu_U^{(2)})$, then

$$\sin^2(\angle(\mathcal{V}, \mathcal{U})) \lesssim \frac{\tilde{W}_{p,k}(\mu^{(1)}, \nu^{(1)}) + \tilde{W}_{p,k}(\mu^{(2)}, \nu^{(2)})}{W_p(\mu^{(1)}, \mu^{(2)})}.$$

Combined with Proposition 8, this implies the claim.

Proposition 14 guarantees the existence of an optimal coupling between $\mu^{(1)}$ and $\mu^{(2)}$ such that if $(X, Y) \sim \gamma$, then $X - Y$ lies in \mathcal{U} almost surely. Let $U \in \mathcal{V}_k(\mathbb{R}^d)$ be such that the column span of U is \mathcal{U} . We have

$$\begin{aligned} W_p(\mu_V^{(1)}, \mu_V^{(2)}) &\leq (\mathbb{E}\|V(X - Y)\|^p)^{1/p} \\ &= (\mathbb{E}\|VU^\top U(X - Y)\|^p)^{1/p} \\ &\leq \|VU^\top\|_{\text{op}} (\mathbb{E}\|X - Y\|^p)^{1/p} = \|VU^\top\|_{\text{op}} W_p(\mu^{(1)}, \mu^{(2)}). \end{aligned}$$

By the definition of V ,

$$\begin{aligned} W_p(\mu^{(1)}, \mu^{(2)}) - W_p(\mu_V^{(1)}, \mu_V^{(2)}) &= W_p(\mu_U^{(1)}, \mu_U^{(2)}) - W_p(\mu_V^{(1)}, \mu_V^{(2)}) \\ &\leq W_p(\mu_U^{(1)}, \mu_U^{(2)}) - W_p(\nu_U^{(1)}, \nu_U^{(2)}) \\ &\quad + W_p(\nu_U^{(1)}, \nu_U^{(2)}) - W_p(\mu_V^{(1)}, \mu_V^{(2)}) \\ &\leq 2 \sup_{U \in \mathcal{V}_k(\mathbb{R}^d)} |W_p(\mu_U^{(1)}, \mu_U^{(2)}) - W_p(\nu_U^{(1)}, \nu_U^{(2)})| \\ &\leq 2\tilde{W}_{p,k}(\mu^{(1)}, \nu^{(1)}) + 2\tilde{W}_{p,k}(\mu^{(2)}, \nu^{(2)}). \end{aligned}$$

Therefore

$$(1 - \|VU^\top\|_{\text{op}}) W_p(\mu^{(1)}, \mu^{(2)}) \leq 2\tilde{W}_{p,k}(\mu^{(1)}, \nu^{(1)}) + 2\tilde{W}_{p,k}(\mu^{(2)}, \nu^{(2)}).$$

Moreover,

$$\begin{aligned} \sin^2(\angle(\mathcal{V}, \mathcal{U})) &= 1 - \sup_{\substack{v \in \mathcal{V}, u \in \mathcal{U} \\ \|u\| = \|v\| = 1}} (u^\top v)^2 \\ &= 1 - \sup_{\substack{x, y \in \mathbb{R}^d \\ \|x\| = \|y\| = 1}} ((Ux)^\top (Vy))^2 \\ &= 1 - \|VU^\top\|_{\text{op}}^2 \\ &\leq 2(1 - \|VU^\top\|_{\text{op}}), \end{aligned}$$

since $\|VU^\top\|_{\text{op}} \leq 1$. The claim follows. \square

ACKNOWLEDGEMENTS

The authors thank Oded Regev for helpful suggestions.

REFERENCES

- [1] Ahidar-Coutrix, A., Gouic, T. L., and Paris, Q. On the rate of convergence of empirical barycentres in metric spaces: curvature, convexity and extendible geodesics. *arXiv:1806.02740*, 2018.
- [2] Alaux, J., Grave, E., Cuturi, M., and Joulin, A. Unsupervised hyperalignment for multilingual word embeddings. *ICLR*, 2019.
- [3] Alvarez-Melis, D., Jaakkola, T. S., and Jegelka, S. Structured optimal transport. *AISTATS*, 2018.
- [4] Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. *ICML*, 2017.
- [5] Bandeira, A. S., Perry, A., and Wein, A. S. Notes on computational-to-statistical gaps: predictions using statistical physics. *arXiv:1803.11132*, 2018.
- [6] Berthet, Q. and Rigollet, P. Optimal detection of sparse principal components in high dimension. *Ann. Statist.*, 41(4):1780–1815, 2013.
- [7] Bigot, J., Cazelles, E., and Papadakis, N. Central limit theorems for Sinkhorn divergence between probability distributions on finite spaces and statistical applications. *arXiv:1711.08947*, 2017.
- [8] Blum, A., Dwork, C., McSherry, F., and Nissim, K. Practical privacy: the sulq framework. *PODS*, 2005.
- [9] Bobkov, S. and Ledoux, M. One-dimensional empirical measures, order statistics and kantorovich transport distances. *preprint*, 2014.
- [10] Bobkov, S. G. and Götze, F. Exponential integrability and transportation cost related to logarithmic Sobolev inequalities. *J. Funct. Anal.*, 163(1): 1–28, 1999.
- [11] Boissard, E. and Le Gouic, T. On the mean speed of convergence of empirical and occupation measures in Wasserstein distance. *Ann. Inst. Henri Poincaré Probab. Stat.*, 50(2):539–563, 2014.
- [12] Bolley, F., Guillin, A., and Villani, C. Quantitative concentration inequalities for empirical measures on non-compact spaces. *Probab. Theory Related Fields*, 137(3-4):541–593, 2007.
- [13] Brennan, M., Bresler, G., and Huleihel, W. Reducibility and computational lower bounds for problems with planted sparse structure. *COLT*, 2018.
- [14] Bubeck, S., Lee, Y. T., Price, E., and Razenshteyn, I. P. Adversarial examples from computational constraints. *ICML*, 2019.
- [15] Cai, T., Ma, Z., and Wu, Y. Optimal estimation and rank detection for sparse spiked covariance matrices. *Probab. Theory Related Fields*, 161(3-4): 781–815, 2015.

- [16] Cai, T. T., Liang, T., and Rakhlin, A. Computational and statistical boundaries for submatrix localization in a large noisy matrix. *Ann. Statist.*, 45(4): 1403–1430, 2017.
- [17] Canas, G. and Rosasco, L. Learning probability measures with respect to optimal transport metrics. *NIPS*, 2012.
- [18] Cazelles, E., Seguy, V., Bigot, J., Cuturi, M., and Papadakis, N. Geodesic PCA versus log-pca of histograms in the wasserstein space. *SIAM J. Scientific Computing*, 40(2), 2018.
- [19] Clatici, S. and Solomon, J. Wasserstein coresets for Lipschitz costs. *ArXiv:1805.07412*, 2018.
- [20] Del Barrio, E., Gamboa, F., Gordaliza, P., and Loubes, J.-M. Obtaining fairness using optimal transport theory. *ICML*, 2019.
- [21] del Barrio, E., Gordaliza, P., Lescornel, H., and Loubes, J.-M. Central limit theorem and bootstrap procedure for wasserstein’s variations with an application to structural relationships between distributions. *Journal of Multivariate Analysis*, 169:341 – 362, 2019.
- [22] del Barrio, E., Inouzhe, H., Loubes, J.-M., Matrán, C., and Mayo-Íscar, A. optimalFlow: Optimal-transport approach to flow cytometry gating and population matching. *arXiv:1907.08006*, 2019.
- [23] Dereich, S., Scheutzow, M., and Schottstedt, R. Constructive quantization: approximation by empirical measures. *Ann. Inst. Henri Poincaré Probab. Stat.*, 49(4):1183–1203, 2013.
- [24] Deshpande, I., Hu, Y.-T., Sun, R., Pyrros, A., Siddiqui, N., Koyejo, S., Zhao, Z., Forsyth, D., and Schwing, A. Max-sliced Wasserstein distance and its use for GANs. *arXiv:1904.05877*, 2019.
- [25] Deutsch, F. The angle between subspaces of a Hilbert space. In *Approximation theory, wavelets and applications (Maratea, 1994)*, volume 454 of *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.*, pages 107–130. Kluwer Acad. Publ., Dordrecht, 1995.
- [26] Diakonikolas, I., Kane, D. M., and Stewart, A. Statistical query lower bounds for robust estimation of high-dimensional Gaussians and Gaussian mixtures. *FOCS*, 2017.
- [27] Dixmier, J. Étude sur les variétés et les opérateurs de Julia, avec quelques applications. *Bull. Soc. Math. France*, 77:11–101, 1949.
- [28] Do Ba, K., Nguyen, H. L., Nguyen, H. N., and Rubinfeld, R. Sublinear time algorithms for earth mover’s distance. *Theory Comput. Syst.*, 48(2):428–442, 2011.
- [29] Dudley, R. M. The speed of mean Glivenko-Cantelli convergence. *Ann. Math. Statist.*, 40:40–50, 1969.

- [30] Feldman, V., Grigorescu, E., Reyzin, L., Vempala, S. S., and Xiao, Y. Statistical algorithms and a lower bound for detecting planted cliques. *J. ACM*, 64(2):Art. 8, 37, 2017.
- [31] Feydy, J., Charlier, B., Vialard, F., and Peyré, G. Optimal transport for diffeomorphic registration. *MICCAI*, 2017.
- [32] Flamary, R., Cuturi, M., Courty, N., and Rakotomamonjy, A. Wasserstein discriminant analysis. *Machine Learning*, 107(12):1923–1945, 2018.
- [33] Forrow, A., Hütter, J.-C., Nitzan, M., Rigollet, P., Schiebinger, G., and Weed, J. Statistical optimal transport via factored couplings. *AISTATS*, 2019.
- [34] Fournier, N. and Guillin, A. On the rate of convergence in Wasserstein distance of the empirical measure. *Probab. Theory Related Fields*, 162(3-4): 707–738, 2015.
- [35] Friedman, J. H. and Tukey, J. W. A projection pursuit algorithm for exploratory data analysis. *IEEE Trans. Computers*, 23(9):881–890, 1974.
- [36] Genevay, A., Peyré, G., and Cuturi, M. Learning generative models with sinkhorn divergences. *AISTATS*, 2018.
- [37] Gozlan, N. Characterization of Talagrand’s like transportation-cost inequalities on the real line. *J. Funct. Anal.*, 250(2):400–425, 2007.
- [38] Gozlan, N. A characterization of dimension free concentration in terms of transportation inequalities. *Ann. Probab.*, 37(6):2480–2498, 2009.
- [39] Gozlan, N. and Léonard, C. A large deviation approach to some transportation cost inequalities. *Probab. Theory Related Fields*, 139(1-2):235–283, 2007.
- [40] Gozlan, N. and Léonard, C. Transport inequalities. A survey. *Markov Process. Related Fields*, 16(4):635–736, 2010.
- [41] Grave, E., Joulin, A., and Berthet, Q. Unsupervised alignment of embeddings with wasserstein procrustes. *AISTATS*, 2019.
- [42] Hoeffding, W. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, 58:13–30, 1963.
- [43] Hütter, J.-C. and Rigollet, P. Minimax rates of estimation for smooth optimal transport maps. *arXiv:1905.05828*, 2019.
- [44] Janati, H., Cuturi, M., and Gramfort, A. Wasserstein regularization for sparse multi-task regression. *AISTATS*, 2019.
- [45] Jiao, J., Han, Y., and Weissman, T. Minimax estimation of the L_1 distance. *IEEE Trans. Inform. Theory*, 64(10):6672–6706, 2018.
- [46] Johnson, W. B. and Lindenstrauss, J. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability (New Haven, Conn., 1982)*, volume 26 of *Contemp. Math.*, pages 189–206. Amer. Math. Soc., Providence, RI, 1984.

- [47] Johnstone, I. M. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics*, 29(2):295–327, 2001.
- [48] Kearns, M. Efficient noise-tolerant learning from statistical queries. *J. ACM*, 45(6):983–1006, 1998.
- [49] Klatt, M., Taming, C., and Munk, A. Empirical regularized optimal transport: Statistical theory and applications. *arXiv:1810.09880*, 2018.
- [50] Kolouri, S., Nadjahi, K., Simsekli, U., Badeau, R., and Rohde, G. K. Generalized sliced Wasserstein distances. *arXiv preprint arXiv:1902.00434*, 2019.
- [51] Kroshnin, A., Spokoiny, V., and Suvorikova, A. Statistical inference for Bures-Wasserstein barycenters. *arXiv:1901.00226*, 2019.
- [52] Kruskal, J. B. Toward a practical method which helps uncover the structure of a set of multivariate observations by finding the linear transformation which optimizes a new “index of condensation”. In *Statistical Computation*, pages 427–440. Elsevier, 1969.
- [53] Kruskal, J. B. Linear transformation of multivariate data to reveal clustering. *Multidimensional scaling: theory and applications in the behavioral sciences*, 1:181–191, 1972.
- [54] Lavenant, H., Claici, S., Chien, E., and Solomon, J. Dynamical optimal transport on discrete surfaces. *ACM Trans. Graph.*, 37(6):250:1–250:16, 2018.
- [55] Le Gouic, T., Paris, Q., Rigollet, P., and Stromme, A. J. Fast convergence of empirical barycenters in Alexandrov spaces and the Wasserstein space. *arXiv:1908.00828*, 2019.
- [56] Ledoux, M. *The concentration of measure phenomenon*, volume 89 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2001.
- [57] Lei, J. Convergence and concentration of empirical measures under Wasserstein distance in unbounded functional spaces. *arXiv:1804.10556*, 2018.
- [58] Liang, T. On the Minimax Optimality of Estimating the Wasserstein Metric. *arXiv:1908.10324*, 2019.
- [59] Lim, S., Lee, S.-E., Chang, S., and Ye, J. C. CycleGAN with a Blur Kernel for Deconvolution Microscopy: Optimal Transport Geometry. *arXiv:1908.09414*, 2019.
- [60] Ma, Z. and Wu, Y. Computational barriers in minimax submatrix detection. *Ann. Statist.*, 43(3):1089–1116, 2015.
- [61] Marton, K. Bounding \bar{d} -distance by informational divergence: a method to prove measure concentration. *Ann. Probab.*, 24(2):857–866, 1996.
- [62] Marton, K. A measure concentration inequality for contracting Markov chains. *Geom. Funct. Anal.*, 6(3):556–571, 1996.

- [63] Maurey, B. Some deviation inequalities. *Geom. Funct. Anal.*, 1(2):188–197, 1991.
- [64] Montavon, G., Müller, K., and Cuturi, M. Wasserstein training of restricted boltzmann machines. *NIPS*, pages 3711–3719, 2016.
- [65] Panaretos, V. M. and Zemel, Y. Statistical aspects of wasserstein distances. *Annual Review of Statistics and Its Application*, 6(1):405–431, 2019.
- [66] Paty, F.-P. and Cuturi, M. Subspace robust Wasserstein distances. *ICML*, 2019.
- [67] Peyré, G. and Cuturi, M. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019.
- [68] Pitié, F., Kokaram, A. C., and Dahyot, R. Automated colour grading using colour distribution transfer. *Computer Vision and Image Understanding*, 107(1-2):123–137, 2007.
- [69] Rabin, J., Peyré, G., Delon, J., and Bernot, M. Wasserstein barycenter and its application to texture mixing. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 435–446. Springer, 2011.
- [70] Ramdas, A., Trillos, N. G., and Cuturi, M. On Wasserstein two-sample testing and related families of nonparametric tests. *Entropy*, 19(2):47, 2017.
- [71] Rigollet, P. and Weed, J. Entropic optimal transport is maximum-likelihood deconvolution. *Comptes Rendus Mathématique*, 356(11):1228 – 1235, 2018.
- [72] Rigollet, P. and Weed, J. Uncoupled isotonic regression via minimum Wasserstein deconvolution. *Information and Inference: A Journal of the IMA*, 04 2019.
- [73] Rolet, A., Cuturi, M., and Peyré, G. Fast dictionary learning with a smoothed wasserstein loss. *MIFODS Semester on Learning under complex structure*, 2016.
- [74] Schiebinger, G., Shu, J., Tabaka, M., Cleary, B., Subramanian, V., Solomon, A., Gould, J., Liu, S., Lin, S., Berube, P., Lee, L., Chen, J., Brumbaugh, J., Rigollet, P., Hochedlinger, K., Jaenisch, R., Regev, A., and Lander, E. S. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943, 2019.
- [75] Schmitz, M. A., Heitz, M., Bonneel, N., Mboula, F. M. N., Coeurjolly, D., Cuturi, M., Peyré, G., and Starck, J. Wasserstein dictionary learning: Optimal transport-based unsupervised nonlinear dictionary learning. *SIAM J. Imaging Sciences*, 11(1):643–678, 2018.
- [76] Seguy, V. and Cuturi, M. Principal geodesic analysis for probability measures under the optimal transport metric. *NIPS*, 2015.
- [77] Singh, S. and Póczos, B. Minimax distribution estimation in Wasserstein distance. *arXiv:1802.08855*, 2018.

- [78] Solomon, J., de Goes, F., Peyré, G., Cuturi, M., Butscher, A., Nguyen, A., Du, T., and Guibas, L. Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Trans. Graph.*, 34(4):66:1–66:11, July 2015.
- [79] Solomon, J., Peyré, G., Kim, V. G., and Sra, S. Entropic metric alignment for correspondence problems. *ACM Trans. Graph.*, 35(4):72:1–72:13, 2016.
- [80] Staib, M., Clatici, S., Solomon, J. M., and Jegelka, S. Parallel streaming wasserstein barycenters. *NIPS*, 2017.
- [81] Talagrand, M. Transportation cost for Gaussian and other product measures. *Geom. Funct. Anal.*, 6(3):587–600, 1996.
- [82] Tameling, C. and Munk, A. Computational strategies for statistical inference based on empirical optimal transport. In *2018 IEEE Data Science Workshop, DSW 2018, Lausanne, Switzerland, June 4-6, 2018*, pages 175–179, 2018.
- [83] Timan, A. F. *Theory of approximation of functions of a real variable*. Dover Publications, Inc., New York, 1994. Translated from the Russian by J. Berry, Translation edited and with a preface by J. Cossar, Reprint of the 1963 English translation.
- [84] Tsybakov, A. B. *Introduction to nonparametric estimation*. Springer Series in Statistics. Springer, New York, 2009. Revised and extended from the 2004 French original, Translated by Vladimir Zaiats.
- [85] Valiant, G. and Valiant, P. A CLT and tight lower bounds for estimating entropy. *Electronic Colloquium on Computational Complexity (ECCC)*, 17:183, 2010.
- [86] Valiant, G. and Valiant, P. The power of linear estimators. *FOCS*, 2011.
- [87] Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. *Arxiv Preprint*, 11 2010.
- [88] Vershynin, R. *High-dimensional probability*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2018. An introduction with applications in data science, With a foreword by Sara van de Geer.
- [89] Villani, C. *Optimal transport*, volume 338 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2009. Old and new.
- [90] Weed, J. and Bach, F. Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *Bernoulli*, 2018. To appear.
- [91] Weed, J. and Berthet, Q. Estimation of smooth densities in wasserstein distance. *COLT*, 2019.
- [92] Wu, Y. and Yang, P. Chebyshev polynomials, moment matching, and optimal estimation of the unseen. *Ann. Statist.*, 47(2):857–883, 2019.

- [93] Yang, K. D., Damodaran, K., Venkatchalapathy, S., Soylemezoglu, A. C., Shivashankar, G., and Uhler, C. Autoencoder and optimal transport to infer single-cell trajectories of biological processes. *bioRxiv*, page 455469, 2018.
- [94] Zemel, Y. and Panaretos, V. M. Fréchet means and procrustes analysis in wasserstein space. *Bernoulli*, 25(2):932–976, 05 2019.

JONATHAN NILES-WEED
COURANT INSTITUTE OF MATHEMATICAL SCIENCES
NEW YORK UNIVERSITY
251 MERCER STREET
NEW YORK, NY 10012-1185, USA
(jnw@cims.nyu.edu)

PHILIPPE RIGOLLET
DEPARTMENT OF MATHEMATICS
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
77 MASSACHUSETTS AVENUE
CAMBRIDGE, MA 02139-4307, USA
(rigollet@math.mit.edu)