

---

# Intrinsic Sliced Wasserstein Distances for Comparing Collections of Probability Distributions on Manifolds and Graphs

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Collections of probability distributions arise in a variety of statistical applications  
2 ranging from user activity pattern analysis to brain connectomics. In practice these  
3 distributions are represented by histograms over diverse domain types including  
4 finite intervals, circles, cylinders, spheres, other manifolds, and graphs. This  
5 paper introduces an approach for detecting differences between two collections of  
6 histograms over such general domains. We propose the intrinsic slicing construction  
7 that yields a novel class of Wasserstein distances on manifolds and graphs. These  
8 distances are Hilbert embeddable, allowing us to reduce the histogram collection  
9 comparison problem to a more familiar mean testing problem in a Hilbert space. We  
10 provide two testing procedures, one based on resampling and another on combining  
11  $p$ -values from coordinate-wise tests. Our experiments in a variety of data settings  
12 show that the resulting tests are powerful and the  $p$ -values are well-calibrated.  
13 Example applications to user activity patterns and spatial data are provided.

## 14 1 Introduction

15 Distributional data arise in a variety of statistical applications. In practice these are not limited to  
16 distributions over real intervals, but are often defined over manifolds and graphs. For instance, even  
17 in the simplest case of analyzing 24-hour activity patterns by constructing histograms of activity  
18 counts by time, the resulting histograms are really supported on a circle rather than an interval on the  
19 real line. If in addition to the time of activity, the observations come with a real number such as the  
20 intensity of the activity, then we end up with a histogram over a cylindrical domain. Spatial datasets  
21 recorded at some geographic region level are another example: one can build a distribution over the  
22 region adjacency graph by capturing the normalized counts of events in each region. When analyzing  
23 distributions over such general domains it is desirable to rely on methods that take into account the  
24 connectivity and geometry of the underlying domain, respect the distributional nature of the data, and  
25 lead to efficient practical algorithms.

26 In this paper we consider the problem of comparing two collections of distributions, namely testing  
27 for homogeneity—whether all of the distributions come from the same *meta-distribution*. While  
28 conceptually similar to two-sample testing, this is a higher order notion in the sense that our units  
29 of analysis are distributions/histograms. Letting  $\mathcal{P}(\mathcal{X})$  denote the set of Borel probability measures  
30 on a metric space  $\mathcal{X}$ , consider the space  $\mathcal{P}(\mathcal{P}(\mathcal{X}))$ . Let  $P, Q \in \mathcal{P}(\mathcal{P}(\mathcal{X}))$ , and assume that we are  
31 given two collections of probability measures  $\{\mu_i\}_{i=1}^{N_1}$  and  $\{\nu_i\}_{i=1}^{N_2}$  that are drawn from  $P$  and  $Q$ ,  
32  $\mu_i \sim P$  and  $\nu_i \sim Q$  in an independent and identical manner. Our goal is to test whether  $P = Q$  and,  
33 moreover, to be able to conduct such tests for general  $\mathcal{X}$ . Such a test would be useful in numerous  
34 practical situations. For instance, an online retailer may aggregate a customer’s monthly activity

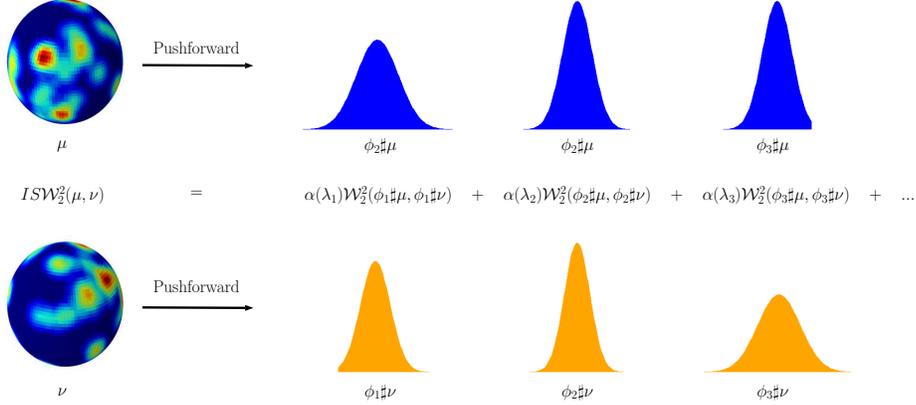


Figure 1: Schematic of the proposed intrinsic slicing construction. Given two probability measures on the sphere (here the darkest blue corresponds to zero mass), different aspects of their dissimilarities become apparent after pushforward to the real line using the eigenfunctions of the Laplace-Beltrami operator,  $\{\phi_i\}$ , in this case spherical harmonics. As a particular example of our general construction, the (squared) intrinsic sliced 2-Wasserstein distance  $ISW_2^2(\cdot, \cdot)$  is the weighted sum of the dissimilarities of the corresponding pushforwards of  $\mu$  and  $\nu$  as measured by squared 2-Wasserstein distance  $W_2^2(\cdot, \cdot)$  on the real line.

35 into a histogram over a cylinder capturing the time of the day and amount of purchase for each  
 36 transaction. By considering collections of histograms for various customer segments, one can conduct  
 37 tests to determine if there are statistically significant differences between behavioral patterns of these  
 38 segments.

39 We attack this problem using insights from recent developments that utilize *Hilbert embeddings* for  
 40 simplifying distributional data problems (see e.g. [18, 24] for particular examples). The simplification  
 41 comes as a result of linearity of Hilbert spaces, which allows adapting existing statistical approaches  
 42 such as functional data methodology to distributional data. A crucial requirement on the embedding  
 43 is that the distance in the embedding space should give a meaningful distance between measures; it is  
 44 this property that renders quantities computed in the embedding space such as means and variances  
 45 meaningful. Thus, embedding constructions should be driven by specifying appropriate distances  
 46 on the space of measures. Of course, not every distance can be embedded and Hilbert embeddable  
 47 distances are called Hilbertian; see [19] for an overview of this notion.

48 The focus in this paper will be on transportation based distances between distributions/histograms  
 49 [25]. Other approaches such as bin-wise treatment of histograms may result in increased variability  
 50 when horizontal variation is present, leading eventually to less powerful methods. Transportation  
 51 based distances are more efficient at capturing this and other aspects of distributional data [3, 17,  
 52 19]. However, adopting the transportation theoretic approaches to our problem immediately hits  
 53 a roadblock beyond the real line case: while 2-Wasserstein distance on the real line is Hilbert  
 54 embeddable, it fails to be so on general domains [19]. The general Hilbert embedding framework of  
 55 [18] is not tied to a distance between probability distributions and so can be problematic for capturing  
 56 the location and variability aspects of distribution collections. In addition, [18] has difficulties in  
 57 higher dimensions and does not provide constructions suitable for manifolds or graphs.

58 Inspired by the sliced 2-Wasserstein distances in high dimensional spaces [11, 12], we introduce  
 59 a new slicing construction (Figure 1) that leverages the eigenvalues and eigenfunctions of the  
 60 Laplace-Beltrami operator on manifolds and graph Laplacians to capture the intrinsic geometry and  
 61 connectivity of the domain. We apply this slicing construction to obtain a novel class of intrinsic sliced  
 62 2-Wasserstein distances on manifolds and graphs. The resulting distances are Hilbert embeddable,  
 63 have a number of desirable properties, and can be truncated to obtain finite-dimensional embeddings.  
 64 Using the corresponding embedding allows us to reduce the histogram collection comparison problem  
 65 to the comparison of means in a high-dimensional Euclidean space. We provide two approaches  
 66 for hypothesis testing and verify via extensive experiments on synthetic and real data examples in a  
 67 variety of data settings that these tests are powerful, and the  $p$ -values are well-calibrated.

68 Comparing with closely related work, while the Generalized Sliced Wasserstein (GSW) distance [11]  
69 sets up the idea of approximating Wasserstein distances using multiple nonlinear projections, it is  
70 presented in extrinsic terms (i.e. Euclidean space) and can suffer from the curse of dimensionality  
71 when a low dimensional data manifold lives in a high-dimensional space. Our choice of eigenfunctions  
72 for projection is very different from the one-parameter function families in GSW. Moreover, the GSW  
73 construction does not directly apply to graphs. While the tree-sliced variant of GSW [14] can be  
74 applied in an intrinsic manner (the clustering variant), it relies on a different type of distance, in the  
75 limit related to the euclidean/geodesic distance. This can be seen by comparing our lower bound to  
76 theirs: our lower bound for ISW is in terms of the MMD using the spectral distance (Proposition 5).  
77 Finally, the robust sliced Wasserstein distance of [13] does make use of the geometric properties of  
78 the underlying manifold. However, their goal is to compute a correspondence between two manifolds  
79 by mapping them into  $\mathbb{R}^d$  using eigenmaps and treating the mapped manifolds as measures in  $\mathbb{R}^d$  and  
80 minimizing some version of Euclidean slicing.

## 81 2 Preliminaries

82 Given a compact metric space  $\mathcal{X}$ , let  $\mathcal{P}(\mathcal{X})$  denote the set of Borel probability measures on  $\mathcal{X}$ . Our  
83 main interest is in the case where  $\mathcal{X}$  is a graph or a manifold with the shortest/geodesic distance as  
84 the metric, and thus the compactness restriction. The 2-Wasserstein distance can be defined on  $\mathcal{P}(\mathcal{X})$   
85 using the metric of  $\mathcal{X}$  as the ground distance [17, 19], giving  $\mathcal{W}_2^{\mathcal{X}} : \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}_{\geq 0}$ ; due  
86 to the repeated use of the real line case we use the shorthand  $\mathcal{W}_2 = \mathcal{W}_2^{\mathbb{R}}$ . Central to our study are  
87 distributions on the space of probability measures  $\mathcal{P}(\mathcal{P}(\mathcal{X})) = (\mathcal{P}(\mathcal{X}), \mathcal{B}(\mathcal{P}(\mathcal{X})))$ , where  $\mathcal{B}(\mathcal{P}(\mathcal{X}))$   
88 is the Borel  $\sigma$ -algebra generated by the topology induced by  $\mathcal{W}_2^{\mathcal{X}}$  [3]. To avoid confusion, we will  
89 refer to the elements of  $\mathcal{P}(\mathcal{P}(\mathcal{X}))$  as *meta-distributions*.

90 Let  $P, Q \in \mathcal{P}(\mathcal{P}(\mathcal{X}))$ , and assume that we are given two collections of probability measures  $\{\mu_i\}_{i=1}^{N_1}$   
91 and  $\{\nu_i\}_{i=1}^{N_2}$  that are drawn from  $P$  and  $Q$ :  $\mu_i \sim P$  and  $\nu_i \sim Q$  in an independent-and-identically-  
92 distributed (hereafter i.i.d.) manner. Our goal is to use this sample to test the null hypothesis  
93 of whether  $P = Q$ . While this is conceptually a two-sample test, note that our data points are  
94 distributions; in practice, the distributions  $\mu_i$  or  $\nu_i$  are given by histograms.

95 *Remark 1.* Let us compare this with the usual two-sample testing. Consider  $P \in \mathcal{P}(\mathcal{P}(\mathcal{X}))$  con-  
96 structed as follows. Let  $\mu^* \in \mathcal{P}(\mathcal{X})$  be a fixed probability measure. Let  $x_1, x_2, \dots, x_A \sim \mu^*$  and  
97 construct the histogram summarizing this sample:  $\frac{1}{A} \sum_{a=1}^A \delta_{x_a}$ . Now,  $\frac{1}{A} \sum_{a=1}^A \delta_{x_a} \in \mathcal{P}(\mathcal{X})$  is one  
98 sample drawn from  $P$ . In our testing scenario one gets the collection  $\{\mu_i\}_{i=1}^{N_1}$ , where each histogram  
99 is obtained as above:  $\mu_i \sim P$ . Similarly, consider  $Q \in \mathcal{P}(\mathcal{P}(\mathcal{X}))$  of the same type based on some  
100 other fixed  $\nu^* \in \mathcal{P}(\mathcal{X})$ , and let  $\{\nu_i\}_{i=1}^{N_2}$  the corresponding collection of histograms. Testing whether  
101  $P = Q$  in the limit boils down to  $\mu^* = \nu^*$ . When compared to the usual two-sample testing this may  
102 seem rather inefficient, requiring  $A$  times more samples (resp.  $N_1 A$  and  $N_2 A$  samples from  $\mu^*$  and  
103  $\nu^*$ ). However, in our setup it is *not assumed* that the histograms in the collections come from meta-  
104 distributions of the above simple type (i.e. all  $\mu_i$  are generated by drawing from the same underlying  
105 distribution  $\mu^*$ ). In fact, the target use-case for our approach is when these histograms are collected  
106 by observing different individuals who have their *person-specific* behaviors/distributions.  $\square$

107 Let  $\mathcal{D}(\cdot, \cdot) : \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}_{\geq 0}$  be a distance between probability distributions.  $\mathcal{D}(\cdot, \cdot)$  is called  
108 *Hilbertian* (this is just a naming convention; no implication that the map is a Hilbert map) if there exist  
109 a Hilbert space  $\mathcal{H}$  and a map  $\eta : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{H}$  such that  $\mathcal{D}(\mu, \nu) = \|\eta(\mu) - \eta(\nu)\|_{\mathcal{H}}$ . For example,  
110 it is well-known that 2-Wasserstein distance on  $\mathcal{X} = \mathbb{R}$  is Hilbertian [19] (also see Section 3.2)  
111 and Maximum Mean Discrepancy (MMD) on any  $\mathcal{X}$  is Hilbertian [8]; however, the 2-Wasserstein  
112 distance  $\mathcal{W}_2^{\mathcal{X}}$  on general  $\mathcal{X}$  is not Hilbertian [19].

113 Since the map  $\eta$  takes every measure on  $\mathcal{X}$  to a point in  $\mathcal{H}$ , we see that a process  $P \in \mathcal{P}(\mathcal{P}(\mathcal{X}))$  gives  
114 a rise to a measure on  $\mathcal{H}$  given by pushforward operation,  $\eta\#P = P \circ \eta^{-1} \in \mathcal{P}(\mathcal{H})$ . In addition, if  
115 a finite dimensional approximation  $\eta_D : \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}^D$  of  $\eta$  is available, then  $\eta_D\#P$  is a measure  
116 on  $\mathbb{R}^D$ . This observation is enormously useful: problems about the elements of the rather abstract  
117 space  $\mathcal{P}(\mathcal{P}(\mathcal{X}))$  are reduced to problems about familiar measures on  $\mathcal{H}$  or even  $\mathbb{R}^D$ . For example,  
118 the usual notions of mean and variance can be applied to the measure  $\eta\#P$  to gain insights about the  
119 meta-distribution  $P$ . The validity of these insights hinges on the  $\eta$ -map coming from a Hilbertian  
120 distance, as distances are central to the statistical quantities of interest.

121 Testing for  $\eta\#P = \eta\#Q$  can serve as a proxy for our original testing problem of  $P = Q$ . As  
 122 typical with two-sample tests, various aspects of the equality  $\eta\#P = \eta\#Q$  can be tested, such as the  
 123 mean or variance equality; unspecific tests of equality can be applied as well. We will concentrate  
 124 on testing certain aspects of the equality so that one can easily drill down on the results. This  
 125 is similar to the regular two-sample testing where checking for equality of, say, means is often  
 126 preferable as it gives immediately interpretable insights, whereas a general test that only says there  
 127 are unspecified differences between the distributions is less useful for interpretation. To obtain succinct  
 128 and interpretable tests we concentrate on the mean of the resulting pushforward measure in  $\mathcal{H}$ .

129 **Definition 1.** For a meta-distribution  $P \in \mathcal{P}(\mathcal{P}(\mathcal{X}))$ , we define its *Hilbert centroid* with respect to  
 130 the Hilbertian distance  $\mathcal{D}$  as  $C_{\eta\#P} = \mathbb{E}_{\mu \sim P}[\eta(\mu)] \in \mathcal{H}$ , assuming it exists.

131 Our testing procedure is based on checking the equality  $C_{\eta\#P} = C_{\eta\#Q}$ , or more explicitly:  
 132  $\mathbb{E}_{\mu \sim P}[\eta(\mu)] = \mathbb{E}_{\nu \sim Q}[\eta(\nu)]$ . Intuitively, each “dimension” of the map  $\eta$  probes some aspect of  
 133 the two involved meta-distributions and makes sure that they are in agreement in expectation. One of  
 134 our testing approaches will use the statistic

$$\mathbb{T}(P, Q) = \|C_{\eta\#P} - C_{\eta\#Q}\|_{\mathcal{H}}^2. \quad (2.1)$$

135 to capture the deviations from equality; this quantity can be written directly in terms of pairwise  
 136 distances.

137 **Proposition 1.** For  $P, Q \in \mathcal{P}(\mathcal{P}(\mathcal{X}))$ , the following holds:

$$\mathbb{T}(P, Q) = \mathbb{E}_{\mu \sim P, \nu \sim Q}[\mathcal{D}^2(\mu, \nu)] - \frac{1}{2} \mathbb{E}_{\mu, \mu' \sim P}[\mathcal{D}^2(\mu, \mu')] - \frac{1}{2} \mathbb{E}_{\nu, \nu' \sim Q}[\mathcal{D}^2(\nu, \nu')].$$

138 Next we give an example of what Hilbert centroid equality means in an important special case.

139 **Example 1.** Let  $\mathcal{X} = [0, T] \subset \mathbb{R}$  with  $\mathcal{D}$  being the 2-Wasserstein distance  $\mathcal{W}_2$ . Given a probability  
 140 measure  $\mu \in \mathcal{P}([0, T])$ , let  $F_\mu$  be its cumulative distribution function:  $F_\mu(x) = \mu([0, x]) = \int_0^x d\mu$ .  
 141 The generalized inverse of cumulative distribution function (CDF) is defined by  $F_\mu^{-1}(s) := \inf\{x \in$   
 142  $[0, T] : F_\mu(x) > s\}$ . The squared 2-Wasserstein distance has a rather simple expression in terms of  
 143 the inverse CDF [19]:

$$\mathcal{W}_2^2(\mu, \nu) = \int_0^1 (F_\mu^{-1}(s) - F_\nu^{-1}(s))^2 ds. \quad (2.2)$$

144 This formula immediately establishes the Hilbertianity of  $\mathcal{W}_2$  through the map  $\eta : \mathcal{P}([0, T]) \rightarrow$   
 145  $L_2([0, T])$  defined by  $\eta(\mu) = F_\mu^{-1}$ . Note that  $\eta$  is invertible for increasing normalized functions  
 146 in the embedding space. Using this insight, we see that the corresponding “average measure” of  
 147  $P \in \mathcal{P}(\mathcal{P}(\mathcal{X}))$  can be introduced via  $P_{\text{av}} = \eta^{-1}(\mathbb{E}_{\mu \sim P}[\eta(\mu)])$ . It is easy to prove that  $P_{\text{av}}$  satisfies  
 148 the following:  $P_{\text{av}} = \arg \min_{\rho \in \mathcal{P}(\mathcal{X})} \mathbb{E}_{\mu \sim P}[\mathcal{W}_2(\mu, \rho)^2]$ , which is the definition of the Fréchet mean,  
 149 see for example [19]. In this setting,  $C_{\eta\#P} = C_{\eta\#Q}$  boils down to having the same Fréchet means,  
 150  $P_{\text{av}} = Q_{\text{av}}$ .  $\square$

151 We will later see that the Hilbert embedding corresponding to the intrinsic sliced 2-Wasserstein  
 152 distance is assembled of embeddings like in Example 1 applied after pushforwards (see Figure 1 for  
 153 an intuition). This means that the resulting equality  $C_{\eta\#P} = C_{\eta\#Q}$  becomes more stringent, making  
 154 it a better proxy for detecting the deviations from  $P = Q$  without losing the interpretability aspect.

### 155 3 Intrinsic Sliced 2-Wasserstein Distance

156 We introduce a Hilbertian version of  $\mathcal{W}_2$  on manifolds and graphs via a construction we call *intrinsic*  
 157 *slicing* due to its use of the domain’s intrinsic geometric properties. To focus our discussion we  
 158 concentrate on the manifold case, as the graph case is simpler and is obtained by replacing the  
 159 Laplace-Beltrami operator by the graph Laplacian.

160 Let  $\lambda_\ell, \phi_\ell; \ell = 0, 1, \dots$  be the eigenvalues and eigenfunctions of the Laplace-Beltrami operator on  
 161  $\mathcal{X}$  with Neumann boundary conditions. The eigenfunctions are sorted by increasing eigenvalue  
 162 and assumed to be orthonormal with respect to some fixed (e.g. uniform) measure on  $\mathcal{X}$ ; also  
 163  $\phi_0 = \text{const}$  and  $\lambda_0 = 0$ . One can define the spectral kernel  $k(x, y) = \sum_\ell \alpha(\lambda_\ell) \phi_\ell(x) \phi_\ell(y)$  and  
 164 the corresponding spectral distance on the manifold  $d(x, y) = k(x, x) + k(y, y) - 2k(x, y) =$   
 165  $\sum \alpha(\lambda_\ell) (\phi_\ell(x) - \phi_\ell(y))^2$ , where  $\alpha : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$  is a function that controls contribution from

166 each spectral band. By setting  $\alpha(\lambda) = e^{-t\lambda}$  for some  $t > 0$ , we get the heat/diffusion kernel and  
 167 the corresponding diffusion distance [4]. Another important case is  $\alpha(\lambda) = 1/\lambda^2$  if  $\lambda > 0$  and  
 168  $\alpha(0) = 0$ , which gives the biharmonic kernel and distance [15]. In both of these constructions  $\alpha(\cdot)$   
 169 is a decreasing function, allowing the smoother low-frequency (i.e. smaller  $\lambda_\ell$ ) eigenfunctions to  
 170 contribute more.

### 171 3.1 Definition and properties

172 A real-valued function  $\phi : \mathcal{X} \rightarrow \mathbb{R}$  can be used to map the manifold  $\mathcal{X}$  onto the real line. Any  
 173 probability measure  $\mu \in \mathcal{P}(\mathcal{X})$  can likewise be projected onto the real line using the pushforward  
 174 of  $\phi$ , which we denote by  $\phi\#\mu = \mu \circ \phi^{-1} \in \mathcal{P}(\mathbb{R})$ . While the pushforward notions used here  
 175 and in previous sections are conceptually the same, for clarity we use  $\#$  for measures and  $\#$  for  
 176 meta-distributions. We define intrinsic slicing as follows.

177 **Definition 2.** Given a function  $\alpha : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$  and a probability distance  $\mathcal{D}(\cdot, \cdot)$  on  $\mathcal{P}(\mathbb{R})$ , we  
 178 define the intrinsic sliced distance  $ISD(\cdot, \cdot)$  on  $\mathcal{P}(\mathcal{X})$  by

$$ISD^2(\mu, \nu) = \sum_{\ell} \alpha(\lambda_{\ell}) \mathcal{D}^2(\phi_{\ell}\#\mu, \phi_{\ell}\#\nu).$$

179 The choice of the Laplacian eigenfunctions in the definition can be justified by a number of their  
 180 properties. Eigenfunctions are intrinsic quantities of a manifold and are ordered by smoothness. Thus,  
 181 they allow capturing the intrinsic connectivity of the underlying domain. Furthermore, due to the  
 182 orthogonality of eigenfunctions, their pushforwards can capture complementary aspects.

183 While the definition is general, our focus in this paper is on the case when  $\mathcal{D} = \mathcal{W}_2$ ; we remind  
 184 that we always use  $\mathcal{W}_2$  to denote the 2-Wasserstein distance on  $\mathcal{P}(\mathbb{R})$ . We call the resulting distance  
 185 *Intrinsic Sliced 2-Wasserstein Distance*, and denote it by  $ISW_2$ . First, we discuss the convergence of  
 186 the infinite sum in Definition 2.

187 **Proposition 2.** If  $\mathcal{X}$  is a smooth compact  $n$ -dimensional manifold and  $\sum_{\ell} \lambda_{\ell}^{(n-1)/2} \alpha(\lambda_{\ell}) < \infty$ ,  
 188 then  $ISW_2$  is well-defined.

189 Next, we prove a number of properties of  $ISD$ .

190 **Proposition 3.** If  $\mathcal{D}$  is a Hilbertian probability distance such that  $ISD$  is well-defined, then (i)  $ISD$   
 191 is Hilbertian, and (ii)  $ISD$  satisfies the following metric properties: non-negativity, symmetry, the  
 192 triangle inequality, and  $ISD(\mu, \mu) = 0$ .

193 *Proof.* By Hilbertian property of  $\mathcal{D}$ , there exists a Hilbert space  $\mathcal{H}^0$  and a map  $\eta^0 : \mathcal{P}(\mathbb{R}) \rightarrow \mathcal{H}^0$   
 194 such that  $\mathcal{D}(\rho_1, \rho_2) = \|\eta^0(\rho_1) - \eta^0(\rho_2)\|_{\mathcal{H}^0}$  for all  $\rho_1, \rho_2 \in \mathcal{P}(\mathbb{R})$ . Plugging this into Definition  
 195 2 we have  $ISD(\mu, \nu) = \|\eta(\mu) - \eta(\nu)\|_{\mathcal{H}}$ , where  $\mathcal{H} = \oplus_{\ell} \mathcal{H}^0$  and the  $\ell$ -th component of  $\eta(\mu)$   
 196 is  $\sqrt{\alpha(\lambda_{\ell})} \eta^0(\phi_{\ell}\#\mu) \in \mathcal{H}$ . The second part of Proposition 3 directly follows from the Hilbert  
 197 property.  $\square$

198 Since  $\mathcal{W}_2$  is Hilbertian on  $\mathcal{P}(\mathbb{R})$ , the application of Proposition 3 yields that  $ISW_2$  is also Hilberitan,  
 199 making it possible to use  $ISW_2$  for our hypothesis tests in Section 4.

200 The following result shows that  $ISW_2$  inherits an important property of the Wasserstein distances,  
 201 namely that the distance between two Dirac delta measures equals to a specific ground distance  
 202 between their locations.

203 **Proposition 4.** When  $\mu = \delta_x(\cdot)$ ,  $\nu = \delta_y(\cdot)$  for two points  $x, y \in \mathcal{X}$ , we have  $ISW_2(\mu, \nu) = d(x, y)$ ,  
 204 where  $d(\cdot, \cdot)$  is the spectral distance corresponding to the choice of  $\alpha(\cdot)$ .

205 For a simple choice of distance  $\mathcal{D}$  on  $\mathcal{P}(\mathbb{R})$ , namely the absolute mean difference, the corresponding  
 206 intrinsic sliced distance is the well-known MMD [8].

207 **Proposition 5.** Let  $\mathcal{D}(\rho_1, \rho_2) = |\mathbb{E}_{x \sim \rho_1}[x] - \mathbb{E}_{y \sim \rho_2}[y]|$  for  $\rho_1, \rho_2 \in \mathcal{P}(\mathbb{R})$ , then the corresponding  
 208  $ISD$  is equivalent to the MMD with the spectral kernel  $k(\cdot, \cdot)$ .

209 When  $k(x, y)$  is the heat kernel, the sliced distance in Proposition 5 is very much like the MMD  
 210 with the Gaussian kernel, with the parameter  $t$  in  $\alpha(\lambda) = e^{-t\lambda}$  controlling the kernel width. Indeed,

211 the two kernels coincide on  $\mathbb{R}^d$ , and on general manifolds Varadhan’s formula gives asymptotic  
 212 equivalence for small  $t$  [2].

213 An interesting insight derived from the above result is that  $ISW_2$  is in a sense a “stronger” distance  
 214 than MMD that uses the corresponding spectral kernel. The  $ISW_2$  compares the quantiles of  
 215 pushforward distributions (Eq. (2.2)), whereas MMD compares their expectations only. We formalize  
 216 this notion next, also providing a theoretical reason for preferring  $ISW_2$  for hypothesis testing.

217 **Proposition 6.**  $MMD(\mu, \nu) \leq ISW_2(\mu, \nu)$  when the same  $\alpha(\cdot)$  is used in both constructions.

218 *Proof.* This follows directly from the fact that for  $\rho_1, \rho_2 \in \mathcal{P}(\mathbb{R})$  the inequality  $|\mathbb{E}_{x \sim \rho_1}[x] -$   
 219  $\mathbb{E}_{y \sim \rho_2}[y]| \leq \mathcal{W}_2(\rho_1, \rho_2)$  holds.  $\square$

220 We are now in a position to prove that  $ISW_2$  is a true metric.

221 **Theorem 1.** If  $\alpha(\lambda) > 0$  for all  $\lambda > 0$ , then  $ISW_2$  is a metric on  $\mathcal{P}(\mathcal{X})$ .

222 We remind that 2-Wasserstein distance can be defined directly on  $\mathcal{P}(\mathcal{X})$  using the geodesic distance  
 223 as the ground metric; we denote this distance as  $\mathcal{W}_2^{\mathcal{X}}$ . Lipschitz properties of the eigenfunctions  
 224 imply the following:

225 **Proposition 7.** There exists a constant  $c$  depending only on  $\mathcal{X} \subseteq \mathbb{R}^n$  such that for all  $\mu, \nu \in \mathcal{P}(\mathcal{X})$   
 226 the inequality  $ISW_2(\mu, \nu) \leq c\mathcal{W}_2^{\mathcal{X}}(\mu, \nu)\sqrt{\sum_{\ell} \lambda_{\ell}^{(n+3)/2} \alpha(\lambda_{\ell})}$  holds.

227 Our final result looks at the quantity  $\mathbb{T}$  defined using  $ISW_2$  by Eq. (2.1). We will be using  $\mathbb{T}$   
 228 computed on finite collections of measures as a test statistic in the next section. We show that it  
 229 enjoys robustness with respect to small perturbations of the measures in the collection.

230 **Proposition 8.** Let  $\{\mu_i\}_{i=1}^N$  and  $\{\nu_i\}_{i=1}^N$  be two collections of probability measures on  $\mathcal{P}(\mathcal{X})$ , such  
 231 that  $\forall i, \mathcal{W}_2^{\mathcal{X}}(\mu_i, \nu_i) \leq \epsilon$ , then  $\mathbb{T}(\{\mu_i\}_{i=1}^N, \{\nu_i\}_{i=1}^N) \leq C^2 \epsilon^2$ . Here  $C = c\sqrt{\sum_{\ell} \lambda_{\ell}^{(n+3)/2} \alpha(\lambda_{\ell})}$  from  
 232 previous proposition and is assumed to be finite.

233 This bound implies that if the distributions in a collection undergo horizontal shifts that are small as  
 234 measured by the geodesic Wasserstein distance  $\mathcal{W}_2^{\mathcal{X}}$ , then  $\mathbb{T}$  is small as well.

### 235 3.2 Approximate Hilbert Embedding

236 An important aspect of  $ISW_2$  is that its Hilbert map  $\eta : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{H}$  can be approximated by a  
 237 finite-dimensional embedding  $\eta_D : \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}^D$  such that  $ISW_2(\mu, \nu) \approx \|\eta_D(\mu) - \eta_D(\nu)\|_{\mathbb{R}^D}$ .  
 238 This is useful for practical computation and for one of our hypothesis testing approaches.

239 Using the formula for  $ISW_2$  on  $\mathcal{P}(\mathbb{R})$  in terms of the quantile function, Eq. (2.2), the Hilbert  
 240 map is defined by  $\eta^0(\mu) = F_{\mu}^{-1}$ . We have  $\mathcal{W}_2(\mu, \nu) = \|\eta^0(\mu) - \eta^0(\nu)\|_{L_2(\mathbb{R})}$ , where the norm  
 241 involves integration. We can discretize the integral using the Riemann sum for equidistant knots  
 242  $s_k = \frac{k-1}{D'}$ ,  $k = 1, \dots, D'$ , define the approximate embedding  $\eta_{D'}^0 : \mathcal{P}(\mathbb{R}) \rightarrow \mathbb{R}^{D'}$  as:

$$\eta_{D'}^0 : \mu \rightarrow \frac{1}{\sqrt{D'}} [F_{\mu}^{-1}(s_1), \dots, F_{\mu}^{-1}(s_{D'})]. \quad (3.1)$$

243 Now,  $\mathcal{W}_2(\mu, \nu) \approx \|\eta_{D'}^0(\mu) - \eta_{D'}^0(\nu)\|_{\mathbb{R}^{D'}}$  with approximation quality depending on the embedding  
 244 dimension  $D'$ .

245 To approximate the Hilbert map for  $ISW_2$  we truncate the series defining  $ISW_2$  and use a finite  
 246 number of eigenfunctions for pushforward:  $\phi_{\ell}$ ,  $\ell = 1, \dots, L$ , where  $\phi_0$  is dropped since it is a constant.  
 247 By inspecting the proof of Proposition 3 and using Eq. (3.1), we can define  $\eta_D : \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}^D$  with  
 248  $D = LD'$  as the concatenation of  $L$  maps:

$$(\eta_D)_{\ell} : \mu \rightarrow \sqrt{\frac{\alpha(\lambda_{\ell})}{D'}} [F_{\phi_{\ell} \# \mu}^{-1}(s_1), \dots, F_{\phi_{\ell} \# \mu}^{-1}(s_{D'})].$$

249 Spectral decompositions of the Laplace-Beltrami operators for general manifolds [4, 20] or graph  
 250 Laplacians can be computed numerically. For applications that involve simple manifolds, the  
 251 eigenvalues and eigenfunctions can be computed analytically (see Appendix).

252 **4 Hypothesis Testing**

253 Let  $\{\mu_i\}_{i=1}^{N_1}$  and  $\{\nu_i\}_{i=1}^{N_2}$  be two i.i.d. collections of measures drawn from  $P, Q \in \mathcal{P}(\mathcal{P}(\mathcal{X}))$   
 254 respectively. Our goal is to use these samples to test the null hypothesis  $H_0 : C_{\eta\#P} = C_{\eta\#Q}$ , where  
 255  $\eta$  is the Hilbert embedding of the sliced distance  $ISW_2$  on  $\mathcal{P}(\mathcal{X})$ .

256 **4.1 Resampling Based Test**

257 We use the quantity  $\mathbb{T}(\cdot, \cdot)$  from Eq. (2.1) as the test statistic. Its sample version is computed by  
 258 replacing the expectations by the empirical means, and excluding the diagonal terms to achieve  
 259 unbiasedness

$$\hat{\mathbb{T}} \equiv \sum_{i,j:i \neq j} \frac{ISW_2^2(\mu_i, \mu_j)}{2N_1(N_1 - 1)} + \sum_{i,j:i \neq j} \frac{ISW_2^2(\nu_i, \nu_j)}{2N_2(N_2 - 1)} - \sum_{i,j} \frac{ISW_2^2(\mu_i, \nu_j)}{N_1 N_2}.$$

260 Note that  $\mathbb{E}\hat{\mathbb{T}} = \mathbb{T}(P, Q)$ . In practice, the  $ISW_2$  values are computed from the approximate  
 261 embedding:  $ISW_2(\rho_1, \rho_2) \approx \|\eta_D(\rho_1) - \eta_D(\rho_2)\|_{\mathbb{R}^D}$ . We denote the resulting statistic by  $\tilde{\mathbb{T}}_{L,D'}$ .

262 The difference between  $\tilde{\mathbb{T}}_{L,D'}$  and the population version (i.e.  $\mathbb{T} - \tilde{\mathbb{T}}_{L,D'}$ ) can be decomposed as  
 263  $(\mathbb{T} - \hat{\mathbb{T}}) + (\hat{\mathbb{T}} - \hat{\mathbb{T}}_L) + (\hat{\mathbb{T}}_L - \tilde{\mathbb{T}}_{L,D'})$ , where the summands inside the terms  $\hat{\mathbb{T}}_L$  and  $\tilde{\mathbb{T}}_{L,D'}$  correspond  
 264 to partial sums that approximate  $ISW_2^2(\cdot, \cdot)$  by  $\sum_{l=1}^L \alpha(\lambda_l) \mathcal{W}_2^2(\phi_l\#, \phi_l\#)$ , and  $\mathcal{W}_2^2(\phi_l\#, \phi_l\#)$  by  
 265  $\|\eta_{D'}(\phi_l\#) - \eta_{D'}(\phi_l\#)\|^2$ , respectively. We show in Appendix that a) summands in the second  
 266 and third terms in the sum can be made infinitesimally small by choosing large enough  $L$  and  $D'$ ,  
 267 respectively; b) an asymptotic result for the first difference can be obtained by extending the tools  
 268 from [8, 23]. These results are based on several assumptions detailed in the Appendix. Combining  
 269 the two results, we establish asymptotic distributions of  $\tilde{\mathbb{T}}_{L,D'}$ :

270 **Theorem 2.** *Assume relevant conditions (see Appendix) hold. Define  $N = N_1 + N_2$ , and suppose*  
 271 *that as  $N_1, N_2 \rightarrow \infty$ , we have  $N_1/N \rightarrow \rho_1, N_2/N \rightarrow \rho_2 = 1 - \rho_1$ , for some fixed  $0 < \rho_1 < 1$ . With*  
 272  *$L \geq L_N, D' \geq D_N$  chosen in an appropriate way (see Appendix), under  $H_0 : C_{\eta\#P} = C_{\eta\#Q}$  we*  
 273 *have*

$$N\tilde{\mathbb{T}}_{L,D'} \rightsquigarrow \sum_{m=1}^{\infty} \gamma_m (A_m^2 - 1),$$

274 where  $A_m \sim N(0, 1)$  for  $m = 1, 2, \dots$ , and  $\gamma_m$  are the eigenvalues of a certain operator that  
 275 depends on  $P$  and  $Q$ . Further, under  $H_1 : C_{\eta\#P} \neq C_{\eta\#Q}$ ,  $\sqrt{N}(\tilde{\mathbb{T}}_{L,D'} - \mathbb{T})$  is asymptotically  
 276 Gaussian with mean 0 and finite variance.

277 We evaluate the power performance of the testing procedure based on  $\tilde{\mathbb{T}}_{L,D'}$  for the sequence of  
 278 contiguous alternatives  $H_{1N} = \{(P, Q) : C_{\mu\#P} = C_{\mu\#Q} + \delta_N, l = 1, 2, \dots\}$ , where the deviation  
 279 from null is quantified collectively by pushforward differences  $\delta_{\ell N} \in \mathcal{H}$ ,  $\delta_N = \oplus_{\ell}(\sqrt{\alpha_{\ell}}\delta_{\ell N})$  that  
 280 are made to approach 0 as  $N \rightarrow \infty$ . The following theorem establishes consistency of our testing  
 281 procedure against a family of such local alternatives.

282 **Theorem 3.** *Assume conditions (i)-(iii) hold, and let  $L, D'$  be chosen as in Theorem 2. Then for*  
 283 *the sequence of contiguous alternatives  $H_{1N}$  such that  $N\|\delta_N\|_{\mathcal{H}^*}^2 \rightarrow \infty$ , the test based on  $\tilde{\mathbb{T}}_{L,D'}$  is*  
 284 *consistent for any  $\alpha \in (0, 1)$ , that is as  $N \rightarrow \infty$  the asymptotic power approaches 1.*

285 **Testing Procedure** In practice, to obtain the  $p$ -value for the  $\tilde{\mathbb{T}}_{L,D'}$ -statistic we use a bootstrap  
 286 procedure. Remember that  $\tilde{\mathbb{T}}_{L,D'}$  is computed via the approximate embedding  $\eta_D$  with  $D = LD'$ .  
 287 The collection  $\{\mu_i\}_{i=1}^{N_1}$  is mapped to the collection  $\{X_i = \eta_D(\mu_i)\}_{i=1}^{N_1}$  of vectors in  $\mathbb{R}^D$  drawn in an  
 288 i.i.d. manner from  $\eta_D\#P = P \circ \eta_D^{-1} \in \mathcal{P}(\mathbb{R}^D)$ . Similarly, for the other collection we have a sample  
 289  $\{Y_i = \eta_D(\nu_i)\}_{i=1}^{N_2}$  drawn from  $\eta_D\#Q$ . Now, the null  $H_0 : C_{\eta\#P} = C_{\eta\#Q}$  implies that the means  
 290 of the distributions  $\eta_D\#P$  and  $\eta_D\#Q$  coincide in  $\mathbb{R}^D$ .

291 The bootstrap null distribution for  $\tilde{\mathbb{T}}_{L,D'}$  can be obtained as follows. Let  $\bar{X}$  and  $\bar{Y}$  be the sample  
 292 means; construct the combined sample  $\{X_i - \bar{X} + \frac{\bar{X} + \bar{Y}}{2}\}_{i=1}^{N_1} \cup \{Y_i - \bar{Y} + \frac{\bar{X} + \bar{Y}}{2}\}_{i=1}^{N_2}$ . This centers both  
 293 samples at  $\frac{\bar{X} + \bar{Y}}{2}$ . Now, from the combined sample we select with replacement  $N_1$  (resp.  $N_2$ ) samples

294 to make bootstrap sample  $\{X_i^b\}_{i=1}^{N_1}$  (resp.  $\{Y_i^b\}_{i=1}^{N_2}$ ). Repeat this process  $B$  times (we take  $B = 1000$   
 295 in our experiments), and collect the null test statistic values  $\tilde{\mathbb{T}}_{L,D'}^b = \tilde{\mathbb{T}}_{L,D'}(\{X_i^b\}_{i=1}^{N_1}, \{Y_i^b\}_{i=1}^{N_2})$  for  
 296  $b = 1, \dots, B$ . The approximate  $p$ -value is then given by:  $p = \frac{1}{B+1} \left( |\{b : \tilde{\mathbb{T}}_{L,D'}^b \geq \tilde{\mathbb{T}}_{L,D'}\}| + 1 \right)$ .

## 297 4.2 Testing via $p$ -value Combination

298 The bootstrap test above incurs a high computational cost and the granularity of the  $p$ -values is  
 299 determined by the number of resamples, which can be too coarse in massive multiple comparison  
 300 settings often seen in industrial applications. Thus, we propose an approach that avoids resampling.

301 As explained above, testing  $H_0 : C_{\eta\#P} = C_{\eta\#Q}$  can be interpreted as testing whether the means  
 302 of the distributions  $\eta_D\#P$  and  $\eta_D\#Q$  coincide in  $\mathbb{R}^D$ . To this end, we adopt the approach pro-  
 303 posed by [22] in a spatial statistics context. First, we apply the Behrens-Fisher-Welch  $t$ -test (with-  
 304 out assuming equality of variances) to each coordinate of the samples  $\{X_i = \eta_D(\mu_i)\}_{i=1}^{N_1}$  and  
 305  $\{Y_i = \eta_D(\nu_i)\}_{i=1}^{N_2}$  to obtain the  $p$ -values  $p_k, k = 1, 2, \dots, D$ . Second, an overall  $p$ -value is com-  
 306 puted via the harmonic mean  $p$ -value combination method which is robust to dependencies [6, 26]:  
 307  $p^H = H \left( D / \left( \frac{1}{p_1} + \frac{1}{p_2} + \dots + \frac{1}{p_D} \right) \right)$ , where the function  $H$  has a known form described in [26].

308 Another approach for combining  $p$ -values is the Cauchy combination test [16], but in our numerical  
 309 experiments we found that the Cauchy combination approach encounters problems when any of the  
 310  $p$ -values is very close to 1, which can happen in our setting due to the form of the embedding  $\eta_D$ .  
 311 Therefore, in contrast to [22], for us the harmonic combination is the only appropriate choice.

312 To guarantee size control, we establish a version of Theorem 1 from [16] for the harmonic mean  
 313  $p$ -value. Assume that a test statistic  $Z \in \mathbb{R}^D$  has null distribution with zero mean and every pair of  
 314 coordinates of  $Z$  follows bivariate Gaussian distribution. Compute the coordinate-wise two-sided  
 315  $p$ -values  $p_k = 2(1 - \Phi(|Z_k|))$  where  $\Phi$  is the standard Gaussian CDF.

316 **Theorem 4.** Let  $p_k, k = 1, \dots, D$  be the null  $p$ -values as above and  $p^H$  computed via harmonic mean  
 317 approach, then

$$\lim_{\alpha \rightarrow 0} \frac{\text{Prob}\{p^H \leq \alpha\}}{\alpha} = 1.$$

318 In the Appendix we show that this theorem applies in our setting, so the proposed procedure  
 319 asymptotically controls the size of the test for small  $\alpha$ . Our experimental results show that the control  
 320 is already achieved for moderate sample sizes and the commonly used  $\alpha = 0.05$ .

## 321 5 Experiments

322 **Synthetic Experiments** We compare the performance of our tests on distributions over finite  
 323 interval, circle, and cylinder with existing methods, and settings of the embedding parameters  $L, D'$ .  
 324 For evaluation, we use empirical power at different degrees of departure from the null hypothesis  
 325 (captured by  $\delta$ ); further details can be found in the Appendix. The summary results presented in  
 326 Figure 2 show that all methods maintain nominal size (power at  $\delta = 0$  is close to  $\alpha$ ). On the inite  
 327 interval: 1) from Figure 2 (a) our combination test (ISD comb) outperformed all the other tests,  
 328 but the bootstrap test (ISD T boot) performs worse than others except Fmaxb; 2) from Figure 2 (b)  
 329 combination test with sliced  $ISW_2$  improves over the unsliced version—with more eigenfunctions, the  
 330 power first improves considerably, then become similar to the unsliced version. For circular domain,  
 331 Figure 2 (c) shows that our tests maintain considerably higher power than existing methods for all  $\delta$ .  
 332 Figure 2 (d) shows that our combination test maintains nominal size on cylindrical domain.

333 **NHANES data on physical activity monitoring** This data [10] contains physical activity pattern  
 334 readings for 6839 individuals. Data for each individual corresponds to activity monitor intensity values  
 335 for 7 days. Since the time dimension is periodic, we get person-specific probability distributions  
 336 over the cylinder  $S^1(T_1) \times [0, T_2)$ . We check if activity patterns vary across age groups. The  
 337  $p$ -value combination test results are shown *below the diagonal* in Table 1. Our method detects  
 338 statistically significant differences between all pairs of groups, except the 36–45 and 46–55 groups.  
 339 As expected, the *control*  $p$ -values—obtained by mixing samples between two age groups and splitting  
 340 arbitrarily—do not concentrate near zero. More details are in the Appendix.

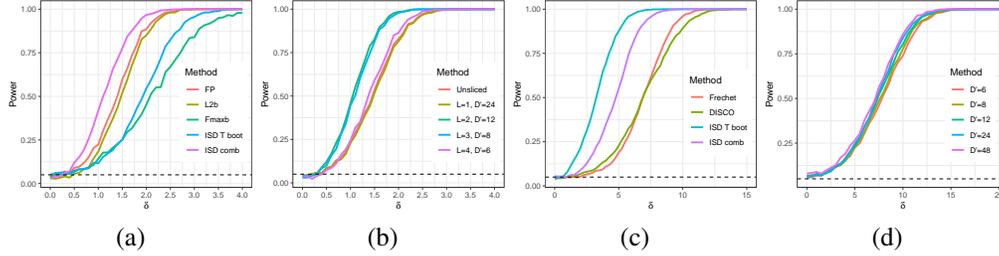


Figure 2: Performance on synthetic finite interval and manifold data. Finite interval: (a) comparison with existing methods—a test based on basis function representation (FP) [7], a sum-type  $\ell_2$  norm-based test (L2b) [27], and a max-type test [28] that uses the maximum of coordinate-wise  $F$  statistic (Fmaxb); (b) unsliced vs. different settings of  $(L, D')$ . Manifold data: (c) circular data, comparing with Fréchet ANOVA [5], and the DISCO nonparametric test [21]; (d) harmonic combination tests on cylindrical data for  $L = 4$ . Dotted lines indicates nominal size of all tests ( $\alpha = 0.05$ ).

Ages	6–15	16–25	26–35	36–45	46–55	56–65	66–75	76–85	Crime Type	Tue vs Thu	Tue vs Sat
6–15		0.394	0.098	0.555	0.882	0.985	0.919	0.997	Theft	0.428	<b>4.2e-06</b>
16–25	<b>1.2e-13</b>		0.575	0.967	0.126	0.921	0.911	0.977	Decept Pract	0.313	<b>0.001</b>
26–35	<b>3.1e-21</b>	<b>2.7e-04</b>		0.459	0.197	0.996	0.919	0.565	Battery	0.430	<b>0.001</b>
36–45	<b>6.1e-22</b>	<b>7.9e-08</b>	<b>0.042</b>		0.864	0.637	0.849	0.991	Robbery	0.119	<b>0.003</b>
46–55	<b>8.2e-22</b>	<b>4.7e-05</b>	<b>0.011</b>	0.343		0.841	0.165	0.554	Narcotics	0.854	<b>0.004</b>
56–65	<b>1.3e-25</b>	<b>0.001</b>	<b>0.001</b>	<b>5.6e-05</b>	<b>0.003</b>		0.991	0.962	Criminal Dam	0.855	<b>0.02</b>
66–75	<b>3.6e-35</b>	<b>7.8e-12</b>	<b>1.5e-11</b>	<b>4.6e-15</b>	<b>1.8e-13</b>	<b>0.001</b>		0.989	Other Offense	0.931	0.052
76–85	<b>3.8e-46</b>	<b>1.4e-26</b>	<b>1.7e-30</b>	<b>8.4e-37</b>	<b>2.1e-35</b>	<b>1.3e-17</b>	<b>6.5e-09</b>		Burglary	0.142	0.261
									Assault	0.997	0.38
									Mot Veh Theft	0.858	0.416

Table 1: Activity intensity comparison across age groups in the NHANES data. Below diagonal:  $p$ -values for to the actual data comparisons. Above diagonal: null  $p$ -values obtained by combining and randomly splitting the two groups. Bold entries correspond to rejected hypotheses with the BH procedure at FDR level 0.1.

Table 2: Chicago Crime analysis  $p$ -values. Bold entries correspond to rejected hypotheses with the BH procedure at FDR level 0.1.

341 **Chicago Crime** We use the Chicago Crimes 2018 dataset [1] to demonstrate the use of our  
342 methodology on histograms over graphs. Each beat (geographic area subdivision used by police)  
343 corresponds to a vertex, and two vertices are connected by an edge if the corresponding beats share  
344 a geographic boundary. For each crime type and day, the normalized counts of that crime type for  
345 each beat gives a daily probability distribution over the graph. Our goal is compare the collection of  
346 distributions of, say, theft occurring on Tuesday to those of Thursday and Saturday. The Tuesday  
347 versus Thursday comparison is intended as a null case, as we do not expect to see any differences  
348 between them [22]. We detect statistically significant differences between Tuesday and Saturday  
349 patterns for six categories of crime, and as expected, no differences between Tuesday and Thursday  
350 patterns. See Appendix for more details and for another graph based application in the context of  
351 Brain Connectomics.

## 352 6 Conclusion

353 The construction of  $ISW_2$  provides a novel embedding of probability distributions into a Hilbert  
354 space. This can be used to adapt many inferential methods to general spaces where the existence of  
355 Fréchet means or higher moments are not guaranteed. The  $ISW_2$  can also be useful for machine  
356 learning applications where prediction targets live in a general domain. Given that rigorous Fréchet  
357 mean-based methodology for such problems has only been proposed recently [9], development of  
358 prediction models for manifold-valued data that are free of restrictive assumptions is an attractive  
359 future line of research.

## References

- 360
- 361 [1] Chicago data portal: Crimes - 2018, 2022. URL [https://data.cityofchicago.org/  
362 Public-Safety/Crimes-2018/3i3m-jwuy](https://data.cityofchicago.org/Public-Safety/Crimes-2018/3i3m-jwuy). 5
- 363 [2] M. Berger. *A Panoramic View of Riemannian Geometry*. Springer-Verlag Berlin Heidelberg,  
364 2003. 3.1
- 365 [3] Jérémie Bigot. Statistical data analysis in the wasserstein space. *ESAIM: ProcS*, 68:1–19, 2020.  
366 doi: 10.1051/proc/202068001. URL <https://doi.org/10.1051/proc/202068001>. 1, 2
- 367 [4] Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational  
368 Harmonic Analysis*, 21(1):5 – 30, 2006. ISSN 1063-5203. doi: [https://doi.org/10.1016/  
369 j.acha.2006.04.006](https://doi.org/10.1016/j.acha.2006.04.006). URL [http://www.sciencedirect.com/science/article/pii/  
370 S1063520306000546](http://www.sciencedirect.com/science/article/pii/S1063520306000546). Special Issue: Diffusion Maps and Wavelets. 3, 3.2
- 371 [5] Paromita Dubey and Hans-Georg Müller. Fréchet analysis of variance for random objects.  
372 *Biometrika*, 106(4):803–821, 10 2019. ISSN 0006-3444. doi: 10.1093/biomet/asz052. URL  
373 <https://doi.org/10.1093/biomet/asz052>. 2
- 374 [6] I. J. Good. Significance tests in parallel and in series. *Journal of the American Statistical  
375 Association*, 53(284):799–813, 1958. ISSN 01621459. 4.2
- 376 [7] T. Gorecki and L. Smaga. A Comparison of Tests for the One-Way ANOVA Problem for  
377 Functional Data. *Computational Statistics*, 30:987–1010, 2015. 2
- 378 [8] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander  
379 Smola. A kernel two-sample test. *J. Mach. Learn. Res.*, 13:723–773, March 2012. ISSN  
380 1532-4435. 2, 3.1, 4.1
- 381 [9] Kyunghye Han, Hans-Georg Müller, and Byeong U. Park. Additive functional regression for  
382 densities as responses. *Journal of the American Statistical Association*, 115:997–1010, 2020. 6
- 383 [10] National Health and Nutrition Examination Survey. 2005-2006 Data Documentation, Code-  
384 book, and Frequencies, 2008. URL [https://www.cdc.gov/Nchs/Nhanes/2005-2006/  
385 PAXRAW\\_D.htm](https://www.cdc.gov/Nchs/Nhanes/2005-2006/PAXRAW_D.htm). 5
- 386 [11] Soheil Kolouri, Kimia Nadjahi, Umut Simsekli, Roland Badeau, and Gustavo Rohde. General-  
387 ized sliced wasserstein distances. In H. Wallach, H. Larochelle, A. Beygelzimer, F. Alché-  
388 Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*  
389 32, pages 261–272. Curran Associates, Inc., 2019. URL [http://papers.nips.cc/paper/  
390 8319-generalized-sliced-wasserstein-distances.pdf](http://papers.nips.cc/paper/8319-generalized-sliced-wasserstein-distances.pdf). 1
- 391 [12] Soheil Kolouri, Phillip E. Pope, Charles E. Martin, and Gustavo K. Rohde. Sliced wasserstein  
392 auto-encoders. In *7th International Conference on Learning Representations, ICLR 2019, New  
393 Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL [https://openreview.net/  
394 forum?id=H1xaJn05FQ](https://openreview.net/forum?id=H1xaJn05FQ). 1
- 395 [13] R. Lai and H. Zhao. Multiscale Nonrigid Point Cloud Registration Using Rotation-Invariant  
396 Sliced-Wasserstein Distance via Laplace–Beltrami Eigenmap. In *SIAM J. Imaging Sci.*, vol-  
397 ume 10 of 2, pages 449–483, 2017. 1
- 398 [14] Tam Le, Makoto Yamada, Kenji Fukumizu, and Marco Cuturi. Tree-sliced variants of wasser-  
399 stein distances. In *Advances in Neural Information Processing Systems*, volume 32, 2019.  
400 1
- 401 [15] Yaron Lipman, Raif M. Rustamov, and Thomas A. Funkhouser. Biharmonic distance. *ACM  
402 Trans. Graph.*, 29(3), July 2010. ISSN 0730-0301. doi: 10.1145/1805964.1805971. URL  
403 <https://doi.org/10.1145/1805964.1805971>. 3
- 404 [16] Yaowu Liu and Jun Xie. Cauchy combination test: A powerful test with analytic p-value calcu-  
405 lation under arbitrary dependency structures. *Journal of the American Statistical Association*,  
406 115(529):393–402, 2020. doi: 10.1080/01621459.2018.1554485. 4.2

- 407 [17] Victor M. Panaretos and Yoav Zemel. Statistical aspects of wasserstein distances. *Annual Review of Statistics and Its Application*, 6(1):405–431, 2019. doi: 10.1146/annurev-statistics-030718-104938. 1, 2
- 410 [18] Alexander Petersen and Hans-Georg Müller. Functional data analysis for density functions by transformation to a hilbert space. *Ann. Statist.*, 44(1):183–218, 02 2016. doi: 10.1214/15-AOS1363. URL <https://doi.org/10.1214/15-AOS1363>. 1
- 413 [19] Gabriel Peyré and Marco Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019. ISSN 1935-8237. doi: 10.1561/22000000073. URL <http://dx.doi.org/10.1561/22000000073>. 1, 2, 2, 1, 1
- 417 [20] Martin Reuter, Franz-Erich Wolter, Martha Shenton, and Marc Niethammer. Laplace-Beltrami eigenvalues and topological features of eigenfunctions for statistical shape analysis. *Computer-Aided Design*, 41(10):739–755, 2009. 3.2
- 420 [21] M. L. Rizzo and G. J. Székely. DISCO analysis: A nonparametric extension of analysis of variance. *Ann. Appl. Stat.*, 4:1034–1055, 2010. 2
- 422 [22] Raif M. Rustamov and James T. Klosowski. Kernel mean embedding based hypothesis tests for comparing spatial point patterns. *Spatial Statistics*, 38:100459, 2020. ISSN 2211-6753. doi: <https://doi.org/10.1016/j.spasta.2020.100459>. URL <http://www.sciencedirect.com/science/article/pii/S2211675320300531>. 4.2, 5
- 426 [23] R. J. Serfling. *Approximation theorems of mathematical statistics*. John Wiley & Sons, 2009. 4.1
- 428 [24] Justin Solomon, Raif M. Rustamov, Leonidas J. Guibas, and Adrian Butscher. Wasserstein propagation for semi-supervised learning. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, volume 32 of *JMLR Workshop and Conference Proceedings*, pages 306–314. JMLR.org, 2014. URL <http://proceedings.mlr.press/v32/solomon14.html>. 1
- 433 [25] Cédric Villani. *Topics in optimal transportation*. Graduate studies in mathematics. American mathematical society, Providence, Rhode Island, 2003. ISBN 0-8218-3312-X. 1
- 435 [26] Daniel J. Wilson. The harmonic mean p-value for combining dependent tests. *Proceedings of the National Academy of Sciences*, 116(4):1195–1200, 2019. ISSN 0027-8424. doi: 10.1073/pnas.1814092116. 4.2
- 438 [27] J.-T. Zhang. *Analysis of Variance for Functional Data*. Chapman and Hall/CRC, first edition, 2013. 2
- 440 [28] J.-T. Zhang, M.-Y. Chen, H.-T. Wu, and B. Zhou. A new test for functional one-way ANOVA with applications to ischemic heart screening. *Computational Statistics & Data Analysis*, 132: 3–17, 2019. 2

## 443 Checklist

- 444 1. For all authors...
- 445 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
- 446
- 447 (b) Did you describe the limitations of your work? [Yes]
- 448 (c) Did you discuss any potential negative societal impacts of your work? [N/A] This is a fairly upstream methodological work.
- 449
- 450 (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
- 451
- 452 2. If you are including theoretical results...
- 453 (a) Did you state the full set of assumptions of all theoretical results? [Yes]

- 454 (b) Did you include complete proofs of all theoretical results? [Yes]
- 455 3. If you ran experiments...
- 456 (a) Did you include the code, data, and instructions needed to reproduce the main experi-  
457 mental results (either in the supplemental material or as a URL)? [Yes]
- 458 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they  
459 were chosen)? [Yes]
- 460 (c) Did you report error bars (e.g., with respect to the random seed after running experi-  
461 ments multiple times)? [No] Evaluation metrics are proportions, so they are already  
462 computed using multiple replications.
- 463 (d) Did you include the total amount of compute and the type of resources used (e.g.,  
464 type of GPUs, internal cluster, or cloud provider)? [No] The computational load for  
465 numerical experiments was minimal enough to not warrant this consideration, i.e.  
466 experiments finish within minutes to hours.
- 467 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 468 (a) If your work uses existing assets, did you cite the creators? [Yes]
- 469 (b) Did you mention the license of the assets? [Yes]
- 470 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
- 471 (d) Did you discuss whether and how consent was obtained from people whose data you're  
472 using/curating? [Yes]
- 473 (e) Did you discuss whether the data you are using/curating contains personally identifi-  
474 able information or offensive content? [Yes] Neither of the real datasets contain any  
475 personally identifiable information or offensive content.
- 476 5. If you used crowdsourcing or conducted research with human subjects...
- 477 (a) Did you include the full text of instructions given to participants and screenshots, if  
478 applicable? [N/A]
- 479 (b) Did you describe any potential participant risks, with links to Institutional Review  
480 Board (IRB) approvals, if applicable? [N/A] IRB approvals have been obtained by  
481 original curators of data.
- 482 (c) Did you include the estimated hourly wage paid to participants and the total amount  
483 spent on participant compensation? [N/A]