Gromov–Wasserstein Transfer Operators[⋆]

Florian Beier[0000−0002−1692−9135], `f.beier@tu-berlin.de`

Institute of Mathematics, Technische Universität Berlin, Straße des 17. Juni 136, 10623 Berlin, Germany

Abstract. Gromov–Wasserstein (GW) transport is inherently invariant under isometric transformations of the data. Having this property in mind, we propose to estimate dynamical systems by transfer operators derived from GW transport plans, when merely the initial and final states are known. We focus on entropy regularized GW transport, which allows to utilize the fast Sinkhorn algorithm and a spectral clustering procedure to extract coherent structures. Moreover, the GW framework provides a natural quantitative assessment on the shape-coherence of the extracted structures. We discuss fused and unbalanced variants of GW transport for labelled and noisy data, respectively. Our models are verified by three numerical examples of dynamical systems with governing rotational forces.

Keywords: optimal transport · Gromov–Wasserstein transport · Perron–Frobenius transfer operators · dynamical systems · coherent structures

1 Introduction

Optimal transport (OT) aims to find an optimal mass transport between two input (marginal) measures according to an underlying cost function. To improve the speed of the numerical computation, Cuturi [9] introduced a regularized OT version which can be solved by the fast and parallelizable Sinkhorn algorithm. Further effort has been made to generalize the OT for different settings as, e.g., unbalanced optimal transport [18], which relaxes the hard matching of the marginal measures. Another line of work pioneered by Mémoli [19] focuses on so-called Gromov–Wasserstein (GW) distances. Here, the inputs have additional structure in the sense of intrinsic (dis-)similarities. The difference to OT is that a meaningful cost function on the product space of the inputs might not be available. Instead, the mass is transported so that pairwise (dis-)similarities are preserved. GW distances are invariant under isometric transformations, making them a valuable tool for e.g. shape classification [4], word alignment [1] or graph matching [29]. For certain applications, a transport which simultaneously takes structural data in the GW sense as well as labelled data in the OT sense into account, is desirable. This is possible in the framework of fused GW transport [26]. Moreover, (fused) GW transport allows for a similar entropic regularization and

unbalanced relaxation as OT [5,23,25]. In [2] the authors propose a framework which extends OT to be invariant to various classes of linear transformations such as e.g. orthogonal transformations. Compared to GW this method is numerically more appealing but has the drawback that it requires the inputs to be embedded in a common space and centered.

Recently, Koltai et al. [16] examined OT-based estimations of dynamical systems from observed initial and final states. More precisely, the authors leveraged solutions to regularized (unbalanced) optimal transport to estimate so-called transfer operators. These are linear operators that characterize dynamical systems in the form of density flows. Furthermore, a clustering procedure based on the spectral information of the estimator was used to extract so-called coherent structures of the dynamical system. Although no unified definition of such structures is available, it is understood that they are persistent in time and space. Coherent structures are of particular interest e.g. in fluid dynamics, since they capture important flow dynamics. This makes precise knowledge of the formation of coherent structures very appealing, since it may lead to a deeper understanding of the dynamics or computational advancements. In [15], the authors assumed instead that the exact transfer operator is known on a finite subset of the full state space. Then, using regularized OT, a finite-dimensional approximation is constructed which limit is a regularized version of the ground truth and exhibits desirable properties, such as retention of the spectral information.

In this paper, we build on the work in [16], but use entropic GW transport plans for constructing transfer operators. This is motivated by the fact that GW transport is readily able to detect isometric transformations such as rotation. Additionally, data labels can be incorporated. We will see that our proposed model includes a quantitative assessment of shape-coherence of the extracted structures.

Outline of the paper. In Section 2, we briefly recall regularized (unbalanced) OT, associated transfer operators and related spectral clustering procedures. GW transport and its (unbalanced) regularized and fused variants are introduced in Section 3. Then, we expand the derivation of transfer operators and spectral clustering towards GW transport plans. In Section 4, we present numerical examples which indicate the potential of our method.

2   Optimal Transport and Transfer Operators

We consider (unbalanced) entropic OT, show how transfer operators can be derived from OT plans, and elaborate on spectral clustering. The derivation of transfer operators will be generalized to GW plans in the next section.

Optimal transport. Let $X, Y \subset \mathbb{R}^d$ be compact sets equipped with the Euclidean distance d_{E}. By $\mathcal{M}^+(X)$, we denote the set of non-negative (Borel) measures and by $\mathcal{P}(X) \subset \mathcal{M}^+(X)$ the set of probability measures on X. Furthermore, let $L^2_\mu(X)$ be the Hilbert space of (equivalence classes) of square integrable functions with respect to the finite measure $\mu \in \mathcal{M}^+(X)$ equipped with the inner product $\langle \cdot, \cdot \rangle_\mu$. By 1_A, we denote the characteristic function on A. For

$\mu, \nu \in \mathcal{M}^+(X)$, the *Kullback–Leibler divergence* is defined by

$$\mathrm{KL}(\mu, \nu) := \int_X \log\left(\tfrac{\mathrm{d}\mu}{\mathrm{d}\nu}\right) \mathrm{d}\mu + \nu(X) - \mu(X),$$

if the Radon–Nikodym derivative $\frac{\mathrm{d}\mu}{\mathrm{d}\nu}$ exists, and by $\mathrm{KL}(\mu, \nu) := \infty$ otherwise. For $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$, a lower semi-continuous cost function $c : X \times Y \to [0, \infty)$ and $\varepsilon > 0$, the *regularized OT problem* is given by

$$\mathrm{OT}_\varepsilon(\mu, \nu) := \min_{\pi \in \Pi(\mu, \nu)} \underbrace{\int_{X \times Y} c(x, y) \, \mathrm{d}\pi + \varepsilon \, \mathrm{KL}(\pi, \mu \otimes \nu)}_{=: F_\varepsilon^{\mathrm{OT}}(\pi)}, \tag{1}$$

where $\Pi(\mu, \nu) := \{\pi \in \mathcal{P}(X \times Y) : P_{1\#}\pi = \mu, P_{2\#}\pi = \nu\}$ with $P_i(x_1, x_2) := x_i$ and push forward measures $P_{i\#}\pi = \pi \circ P_i^{-1}$, $i = 1, 2$. Elements of $\Pi(\mu, \nu)$ are called *transport plans*. For $\varepsilon = 0$, we obtain the unregularized optimal transport $\mathrm{OT}(\mu, \nu)$. The minimizer in (1) is called (entropic) optimal transport plan $\hat{\pi}_\varepsilon$. In the following, we will mainly use $c = d_{\mathrm{E}}^2$ which leads to the *Wasserstein distance* $\mathrm{OT}(\mu, \nu)^{\frac{1}{2}}$. The dual problem of OT_ε is

$$\mathrm{OT}_\varepsilon(\mu, \nu) = \max_{(f, g) \in L_\mu^\infty(X) \times L_\nu^\infty(Y)} \left\{ \int_X f \, \mathrm{d}\mu + \int_Y g \, \mathrm{d}\nu \right.$$
$$\left. - \varepsilon \int_{X \times Y} \exp\left(\frac{f(x) + g(y) - c(x, y)}{\varepsilon}\right) - 1 \, \mathrm{d}(\mu \otimes \nu) \right\}$$

Optimal potentials $\hat{f}_\varepsilon \in L_\mu^\infty(X)$, $\hat{g}_\varepsilon \in L_\nu^\infty(Y)$ exist and are unique on $\mathrm{supp}(\mu)$ and $\mathrm{supp}(\nu)$ up to an additive constant. They are related to $\hat{\pi}_\varepsilon$ by

$$\hat{\pi}_\varepsilon = \exp\left(\frac{\hat{f}_\varepsilon(x) + \hat{g}_\varepsilon(y) - c(x, y)}{\varepsilon}\right) (\mu \otimes \nu) =: k_\varepsilon(\mu \otimes \nu). \tag{2}$$

For atomic measures, the solution can be approximated efficiently by Sinkhorn's algorithm. In some applications, it is useful to deal with *regularized unbalanced OT*

$$\mathrm{UOT}_{\varepsilon, \kappa}(\mu, \nu) := \min_{\pi \in \mathcal{M}^+(X \times Y)} F_\varepsilon^{\mathrm{OT}}(\pi) + \kappa \left(\mathrm{KL}(P_{1\#}\pi, \mu) + \mathrm{KL}(P_{2\#}\pi, \nu)\right), \ \kappa > 0,$$

which relaxes the hard marginal constraints on the objective to penalizing the KL divergence of its marginals with respect to the inputs. Unbalanced optimal transport is treated in detail in [18] and its regularized version in [22]. Similarly as in the balanced case, there is a dual problem formulation with optimal potentials $(\hat{f}_{\varepsilon, \kappa}, \hat{g}_{\varepsilon, \kappa}) \in L_\mu^\infty(X) \times L_\nu^\infty(Y)$ and the optimal transport plan is given by

$$\hat{\pi}_{\varepsilon, \kappa} = \exp\left(\frac{\hat{f}_{\varepsilon, \kappa}(x) + \hat{g}_{\varepsilon, \kappa}(y) - c(x, y)}{\varepsilon}\right) (\mu \otimes \nu) =: k_{\varepsilon, \kappa}(\mu \otimes \nu).$$

A generalization of Sinkhorn's algorithm can be used to solve the corresponding discrete problem, see [22].

Transfer operators. *Transfer operators*, also known as *Perron-Frobenius operators*, are linear operators which characterize dynamical systems in the form of density flows [14]. We consider transfer operators derived from entropic transport plans as in [16]. Here we restrict ourselves to the balanced setting as the unbalanced case follows in a similar way. To this end, we associate $\hat{\pi}_\varepsilon$ in (2) with the transfer operator $K_\varepsilon : L^2_\mu(X) \to L^2_\nu(Y)$ given by

$$(K_\varepsilon \psi)(y) := \int_X k_\varepsilon(x,y)\psi(x)\,\mathrm{d}\mu(x).$$

Figuratively, K_ε captures the structure of the transport of $\hat{\pi}_\varepsilon$ independent of the marginal masses. Since $\hat{\pi}_\varepsilon \in \Pi(\mu,\nu)$, it holds

$$\int_X k_\varepsilon(x,y)\,\mathrm{d}\mu(x) = 1_Y \;\; \nu\text{-a.e.} \quad \text{and} \quad \int_Y k_\varepsilon(x,y)\,\mathrm{d}\nu(y) = 1_X \;\; \mu\text{-a.e.}.$$

In particular, for atomic measures $\mu = \sum_{i=1}^m \mu(i)\delta_{x_i}$ and $\nu = \sum_{j=1}^n \nu(j)\delta_{y_j}$ and an optimal transport plan $\hat{\pi}_\varepsilon = \sum_{i,j=1}^{m,n} \hat{\pi}_\varepsilon(i,j)\delta_{x_i,y_j}$ using the matrix-vector notation $\mu := (\mu(i))_{i=1}^m$, $\nu := (\nu(j))_{j=1}^n$, $D_\mu := \mathrm{diag}(\mu)$ and $\hat{\pi}_\varepsilon := (\hat{\pi}_\varepsilon(i,j))_{i,j=1}^{m,n}$, the transfer kernel and operator are given by

$$k_\varepsilon = D_\mu^{-1}\,\hat{\pi}_\varepsilon\,D_\nu^{-1} \quad \text{and} \quad K_\varepsilon = D_\nu^{-1}\,\hat{\pi}_\varepsilon^{\mathrm{T}}.$$

Spectral clustering. In order to extract coherent structures in dynamical systems, we can apply a spectral clustering procedure on K_ε, see [16]. The clustering premise is just the knowledge of two observations from $\mu \in \mathcal{P}(X)$ and $\nu \in \mathcal{P}(Y)$ in a dynamical system without any knowledge of the true dynamics. The goal is to find measurable partitions $X = X_1 \dot\cup X_2$, $Y = Y_1 \dot\cup Y_2$ fulfilling ideally

$$K_\varepsilon 1_{X_k} = 1_{Y_k} \quad \text{and} \quad \mu(X_k) = \nu(Y_k), \quad k = 1,2.$$

These conditions may be interpreted as coherence and mass preservation of the partitions. One way to tackle this problem is to consider the following optimization problem

$$\max_{X_1 \dot\cup X_2 = X, Y_1 \dot\cup Y_2 = Y} \left\{ \frac{\langle K_\varepsilon 1_{X_1}, 1_{Y_1}\rangle_\nu}{\mu(X_1)} + \frac{\langle K_\varepsilon 1_{X_2}, 1_{Y_2}\rangle_\nu}{\mu(X_2)} \right\},$$

which is usually relaxed to

$$\max_{(\varphi,\psi)\in L^2_\mu(X)\times L^2_\nu(Y)} \left\{ \frac{\langle K_\varepsilon \varphi, \psi\rangle_\nu}{\|\varphi\|_\mu \|\psi\|_\nu} : \langle \varphi, 1_X\rangle_\mu = \langle \psi, 1_Y\rangle_\nu = 0 \right\}. \tag{3}$$

Since K_ε is bounded and non-negative $(\mu \otimes \nu)$-a.s., it follows that the largest singular value of $K_\varepsilon^* K_\varepsilon$ is simple [14, Lem. 3]. Moreover, the largest singular value of K_ε is 1 and the corresponding left and right singular functions are 1_X and 1_Y, respectively. Notably, $(1_X, 1_Y)$ are not included by the constraints in (3). Hence, a maximizing pair $(\hat{\varphi}, \hat{\psi})$ in (3) is given by the right and left singular functions of K_ε associated to the second largest singular value of K_ε. The desired partitioning is then readily obtained by thresholding $(\hat{\varphi}, \hat{\psi})$ at zero. Solving (3) in practice amounts to computing a (truncated) singular value decomposition of K_ε.

3 Transfer Operators from GW Transport Plans

In both references [15,16], the assumption on the underlying dynamics is that they are compliant with an optimal transport. For certain situations this might not be the case. Consider e.g. particles on the two-dimensional unit disk with a driving rotational force. If the rotation angle between two observations is large, OT will not be able to recover this dynamic, see the first Example in Section 4. A transport setting which naturally handles isometric transforms such as rotation is given by the framework of GW transport [19]. As before, we consider compact state spaces $X, Y \subset \mathbb{R}^d$ and measures $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$. In the contrast to classic OT, a cost function on the product space $X \times Y$ is not required. Instead we seek the preservation of the internal structure of the spaces. Here we focus on the Euclidean metrics, for generalizations see [24]. We set $d_X := d_{\mathrm{E}}|_{X \times X}$. Then the triples $\mathbb{X} := (X, d_X, \mu)$, $\mathbb{Y} := (Y, d_Y, \nu)$ are called *metric measure (mm-) spaces*. We introduce the notation $\mu^\otimes := \mu \otimes \mu$. For $\varepsilon > 0$, the *regularized GW transport* between two mm-spaces $\mathbb{X}$ and $\mathbb{Y}$ is defined by

$$\mathrm{GW}_\varepsilon(\mathbb{X}, \mathbb{Y}) := \inf_{\pi \in \Pi(\mu, \nu)} F_\varepsilon^{\mathrm{GW}}(\pi),$$

$$F_\varepsilon^{\mathrm{GW}}(\pi) := \int_{(X \times Y)^2} (d_X(x, x') - d_Y(y, y'))^2 \, \mathrm{d}\pi(x, y) \, \mathrm{d}\pi(x', y') + \varepsilon \, \mathrm{KL}\big(\pi^\otimes, (\mu \otimes \nu)^\otimes\big).$$

In contrast to OT, we regularize with the quadratic KL divergence as in [23]. For $\varepsilon = 0$, we obtain the unregularized GW transport GW which was originally introduced in [19]. Notably, $\mathrm{GW}^{\frac{1}{2}}$ defines a metric on the space of mm-spaces up to identification by measure-preserving isometries. More precisely, $\mathrm{GW}(\mathbb{X}, \mathbb{Y}) = 0$ if and only if there exists an isometry $I : X \to Y$ with $\nu = I_\#\mu$. In this case, $(\mathrm{id}, I)_\#\mu$ is an optimal GW plan. In particular, this shows the invariance of GW with respect to isometric transformations. Figuratively, optimal GW plans are such that whenever they transport (infinitesemal) mass from x to y and x' to y' one has $d_X(x, x') \approx d_Y(y, y')$ which favors a near-isometric transport.

Similar to OT_ε, GW_ε admits unbalanced versions [23], we focus on marginal penalization using KL. For $\varepsilon, \kappa > 0$, the *unbalanced regularized GW transport* is defined by

$$\mathrm{UGW}_{\varepsilon, \kappa}(\mathbb{X}, \mathbb{Y}) = \inf_{\pi \in \mathcal{M}^+(X \times Y)} F_\varepsilon^{\mathrm{GW}}(\pi) + \kappa \left(\mathrm{KL}((P_{1\#}\pi)^\otimes, \mu^\otimes) + \mathrm{KL}((P_{2\#}\pi)^\otimes, \nu^\otimes) \right).$$

Here the marginals of optimal plans differ from the inputs whenever an exact matching results in large values under the functional $F_\varepsilon^{\mathrm{GW}}$. This can make $\mathrm{UGW}_{\varepsilon, \kappa}$ somewhat robust to outliers.

When working with labelled data, we might be interested in a transport plan which preserves the internal geometrical information in the form of metrics as well as feature information in the form of labels. This leads to a fused version of the GW and the Wasserstein distance. To incorporate label information, we introduce an additional set $A \subset \mathbb{R}^m$ endowed with $d_A := d_{\mathrm{E}}|_{A \times A}$. We assume

that each point in X, Y admits only one label, which we characterize by label functions $l_X : X \to A$, $l_Y : Y \to A$, respectively. Clearly, a more general treatment would be to consider distributions in the label space as in e.g. [26]. In our case, the *regularized fused GW distance* is defined by

$$\mathrm{FGW}_\varepsilon((\mathbb{X}, l_X), (\mathbb{Y}, l_Y)) \coloneqq \inf_{\pi \in \Pi(\mu,\nu)} F_\varepsilon^{\mathrm{GW}}(\pi) + \int_{X \times Y} d_A\big(l_X(x), l_Y(y)\big) \, \mathrm{d}\pi(x, y).$$

As with the original formulation, the marginal constraints may be relaxed in the same way which leads to an unbalanced, fused variant $\mathrm{UFGW}_\varepsilon^\kappa$ which was discussed in [5,25].
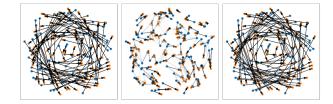
The previously discussed GW formulations are quadratic with respect to the objective plan which renders them numerically challenging. For our numerical experiments below, we rely on a class of simple iterative algorithms which are based on block-coordinate relaxations. The main idea consists of alternately fixing one plan while minimizing with respect to the other. The problem that is then minimized in each iteration step can be written as an entropic OT problem for which Sinkhorn's algorithm can be leveraged. Details regarding this procedure can be found in [21] (balanced GW), [23] (unbalanced GW) and [5,25] (unbalanced, fused GW). Solutions $\hat{\pi}_\varepsilon$ obtained with this procedure are also solutions to an entropic (unbalanced) OT problem and thus have the form (2), i.e. it holds $\hat{\pi}_\varepsilon = k_\varepsilon \, (\mu \otimes \nu)$. Ultimately, this allows us to apply the spectral clustering procedure on the associated transfer operator K_ε as described in Section 2. The next remark highlights another benefit of GW over OT transfer operators for extracting coherent structures.

Remark 1 (Quantitative assessment of shape-coherence). Let $\hat{\pi}_\varepsilon$ be an optimal GW plan between $\mathbb{X}$ and $\mathbb{Y}$ with associated transfer operator K_ε and X_i, Y_i, $i = 1, 2$ the spectral clustering partition. Even if the partitions satisfy $K_\varepsilon 1_{X_i} \approx 1_{Y_i}$ and $\mu(X_i) \approx \nu(Y_i)$, it may be that the intrinsic shapes of X_i and Y_i, $i = 1, 2$ differs significantly. It depends on the application, if these structures should be considered coherent or not. The GW framework readily gives us the possibility for a quantitative assessment of shape-coherence by evaluating the GW functional at $\hat{\pi}_\varepsilon$ restricted to $X_i \times Y_i$, $i = 1, 2$. The closer the evaluation is to 0, the more the associated partitions can be considered shape-coherent or isometric under the transfer operator K_ε. We apply this for Example 3 in Section 4.

4   Numerical Examples

In this section we provide three examples of our proposed GW transfer method. In OT comparisons we use the quadratic Euclidean cost function. We partly rely on the Python Optimal Transport library [12]. For our experiments we aim to set the entropic regularization parameter $\varepsilon > 0$ as small as possible while avoiding numerical overflow.

1. Particles on a rotating disk. First, we are interested in the ability of OT plans to recover the dynamics of a rotating system and compare it with a GW

Fig. 1. An example of X (blue), $R_{\frac{\pi}{2}}(X)$ (orange), as well as the associated transfer kernels k_{true} (left), $k_{\varepsilon}^{\mathrm{W}}$ (middle) and $k_{\varepsilon}^{\mathrm{GW}}$ (right). The arrows are drawn from x to y with opacity proportional to the respective kernel at (x, y).

based approach. We consider $n = 50$ uniformly sampled particles on the 2D unit disk $D := \{x : \|x\|_2 \leq 1\} \subset \mathbb{R}^2$. Let $X \subset D$ be the set of particles. We consider a counter-clockwise rotation of degree $\theta \in (0, 2\pi)$. More precisely, the true transfer operator is characterized by the bijective map $R_\theta : D \to D$ given by

$$(r\cos(\phi), r\sin(\phi)) \mapsto (r\cos(\phi + \theta), r\sin(\phi + \theta)), \quad r \in [0, 1], \theta \in [0, 2\pi).$$

We focus on the transfer associated to kernel $k_{\mathrm{true}}(x, y) = \delta_{\{R_\theta(x)=y\}}$. An illustration of X and R_θ is shown in Figure 1. We investigate how well the GW transfer operator estimates the true transfer operator for $\theta = \frac{\pi}{30}, \frac{2\pi}{30}, \ldots, \pi$. To this end, we sample the initial state X and compute the GW transport plan with $\varepsilon = 0.0008$ between between the uniform distributions on X and $R_\theta(X)$, respectively. As discussed, all plans admit the form (2) for respective kernels $k_{\varepsilon}^{\mathrm{W}}, k_{\varepsilon}^{\mathrm{GW}}$. To compare the performance we consider the error measure

$$\mathrm{e}(k_{\varepsilon}^{\bullet}) := \frac{1}{n^2} \sum_{x \in X} \sum_{y \in R_\theta(X)} k_{\varepsilon}^{\bullet}(x, y) d_{\mathrm{E}}(R_\theta(x), y), \quad \bullet \in \{\mathrm{W}, \mathrm{GW}\}.$$

Intuitively, this gives us the mean Euclidean distance when comparing the transfer operator associated to the kernel against the true transfer.

The right-hand side of Figure 1 shows the qualitative difference between the OT and GW-based approaches for one example with $\theta = \frac{\pi}{2}$. A quantitative comparison is given on the left-hand side of Figure 2. More precisely, we sampled 10 independent choices of X to obtain 10 OT plans $\pi_{\varepsilon}^{\mathrm{W},1}, \ldots, \pi_{\varepsilon}^{\mathrm{W},10}$ and 10 GW plans $\pi_{\varepsilon}^{\mathrm{GW},1}, \ldots, \pi_{\varepsilon}^{\mathrm{GW},10}$ for each angle θ. We plot the mean errors $\frac{1}{10} \sum_{i=1}^{10} \mathrm{e}(k_{\varepsilon}^{\bullet,i})$, $\bullet \in \{\mathrm{W}, \mathrm{GW}\}$ as a function of the angle θ. As expected, for large values of θ, the OT-based transfer operator is a poor estimator. This is evident since, e.g. for a 90 degrees rotation, points are transferred far distances which is sub-optimal in the OT sense. Even for smaller angles such as 18 degrees, we observe a mean error of 0.15. On the other hand, the GW based approach recovers R_θ nearly exactly in all cases. In the previous example, $Y = R_\theta(X)$ was given by an exact rotation of X. However, in practice the observed end state Y of the dynamical system might be a noisy version of $R_\theta(X)$. Hence, we repeat the previous experiment, where this time $Y = R_\theta(X) + m\eta$ with $\eta \sim \mathcal{U}([-0.1, 0.1]^2)$
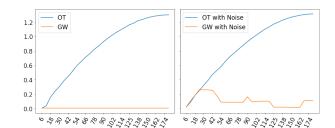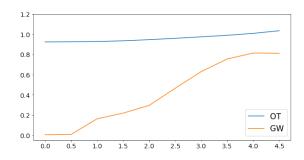
Fig. 2. The mean errors plotted against rotation angle θ in degrees without noise (left) and with noise (right).
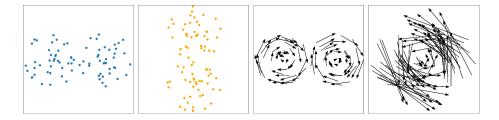

Fig. 3. The mean errors for fixed $\theta = \frac{\pi}{2}$ plotted against the noise magnitude m.

and $m = 1$. To make this comparable to the previous experiment, we consider the same sampled initial states X as above. We proceed as before and plot the error of the OT and GW-based approach on the right-hand side of Figure 2. For small angles, GW remains comparable to OT whereas for large angles a better estimation is achieved by utilizing GW. Finally, we repeat the procedure this time for a fixed angle $\theta = \frac{\pi}{2}$ and for varying $m = 0.5, 1, 1.5, \ldots, 4.5$. The result is plotted in Figure 3.
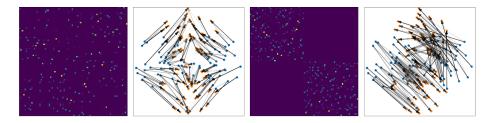
2. Multiple rotating disks. In our next example, let $\theta = \pi/2$, and D, R_θ as above. In addition, for $i = 1, 2$, we consider

$$D_i = \{x \in D : \|x - x^{(i)}\| \le 1/2\}, \quad x^{(i)} = (-1/2, 0), \quad x^{(i)} = (1/2, 0).$$

We set $F = (R^{(1)} + R^{(2)}) \circ R_\theta$, where $R^{(i)}$ constitutes a rotation of $-\pi/4$ around $x^{(i)}$, restricted to D_i, $i = 1, 2$. Let $n = 80$, we uniformly sample $n/2$ points of D_1 and D_2, respectively. Denote the entire set of n points by $X \subset D$. Let $Y = F(X)$ and equip X and Y with the uniform distribution. Figure 4 illustrates X, Y and F. We focus on the estimation of the transfer operator associated to $k_{\text{true}}(x, y) = \delta_{\{F(x) = y\}}$. We compute an OT plan $\pi_\varepsilon^{\text{W}}$ and an GW plan $\pi_\varepsilon^{\text{GW}}$ both with $\varepsilon = 0.001$. Illustrations of the matrices $\pi_\varepsilon^{\text{W}}, \pi_\varepsilon^{\text{GW}}$ as well as a visualizations of the transfer operators $K_\varepsilon^\bullet$, associated to respective kernels $k_\varepsilon^\bullet$, $\bullet \in \{\text{W}, \text{GW}\}$ are provided in Figure 5. Clearly, neither approach is able to recover the ground truth. However, the figure indicates that $K_\varepsilon^{\text{GW}}$ transfers most of the mass from

Fig. 4. Left to right: X, Y, $(R^{(1)} + R^{(2)})$ and R_θ.

Fig. 5. Left to right: Matrix $\pi_\varepsilon^{\mathrm{W}}$, transfer kernel $k_\varepsilon^{\mathrm{W}}$, matrix $\pi_\varepsilon^{\mathrm{GW}}$, transfer kernel $k_\varepsilon^{\mathrm{GW}}$.

D_i to $F(D_i)$, $i = 1, 2$, while the OT-based approach does not. We apply the spectral clustering procedure, i.e. we compute the left and right eigenvectors associated to the second largest eigenvalue of $K_\varepsilon^\bullet$, $\bullet \in \{\mathrm{W}, \mathrm{GW}\}$ and present them in Figure 6. As expected, the partitioning according to $K_\varepsilon^{\mathrm{GW}}$ is able to find both coherent disks. Now, the local dynamics within the partitions can readily be obtained by computing the GW transport of the partitioned subspaces.

We conclude this example by remarking that the correct identification of the discs may also fail and is not stable with respect to noise. This is due to the fact that e.g. $\mathbb{X}$ is almost isometric to a 180 degrees rotation as well as a reflection along the vertical axis. If the inputs are subjected to noise, an optimal GW plan might match D_1 with $F(D_2)$ and D_2 with $F(D_1)$.
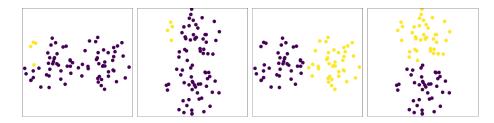
3. Vorticity field of the 2D Navier–Stokes equation. Finally, we consider a two-dimensional flow in time which behaves according to the 2D Navier–Stokes equations on the square $[0, 2\pi]^2$ (periodic boundary conditions)

$$\partial_t u + (u \cdot \nabla)u = -\nabla p + v\nabla^2 u$$
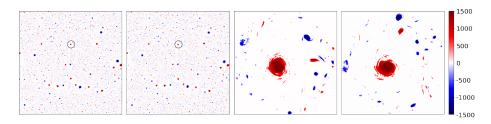$$\nabla \cdot u = 0,$$

where $u : [0, T] \times [0, 2\pi]^2 \to \mathbb{R}^2$ is the velocity, $p : [0, T] \times [0, 2\pi]^2 \to \mathbb{R}^2$ the pressure and $v \in \mathbb{R}$ the kinematic viscosity. Numerically, it is more efficient to solve the scalar advection-diffusion equation

$$\partial_t \omega + (u \cdot \nabla)\omega = v\nabla^2 \omega, \tag{4}$$

where $\omega = \partial_x u_y - \partial_y u_x$ is the vorticity of u. Following [27, Sec IV], the equation is solved in the Fourier domain after a adding a small-scale forcing term and
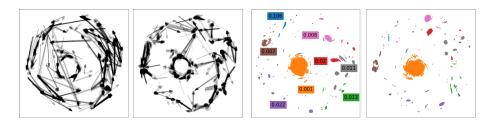
Fig. 6. The spaces X, Y coloured according to the sign of the left and right eigenvector of $K_\varepsilon^{\mathrm{W}}$ (left) and $K_\varepsilon^{\mathrm{GW}}$ (right) corresponding to the second largest eigenvalue.

Fig. 7. Left: Two thresholded snapshots ω_0, ω_1 of a direct numerical simulation of (4) on a 4096x4096 pixel grid. The images on the right-hand side show the marked circular patch for both time-steps.

a large-scale damping function on a 4096x4096 grid. Ultimately, we obtain two time snapshots ω_0, ω_1 of the vorticity field, which we restrict to $|\omega_i| \geq 600$. The snapshots as well as a zoom into a circular patch with a 290 pixel diameter is shown in Figure 7. As we can see, the flow exhibits coherent structures in the form of vortices on large and small scales. Large vortices essentially determine most of the local dynamics. This can be seen for instance in the in selected patch, where smaller vortices are rotating around the large center vortex.

We proceed to estimate the dynamics of the extracted patch. Similarly to our motivating example, we compare the OT and GW transfer operators. From ω_0 and ω_1 we extract the mm-spaces $\mathbb{X} = (X, d_X, \mu)$ and $\mathbb{Y} = (Y, d_Y, \nu)$, respectively. More precisely, X and Y are the sets of patch points in $\mathbb{R}^2$ where $|\omega_0| \geq 600$ and $|\omega_1| \geq 600$. Furthermore, d_X and d_Y are the normalized Euclidean metrics on X and Y, respectively. Finally, μ, ν are the (fully supported) probability measures proportional to the absolute value of the vorticity field. For our model we want to prohibit the transport between positive and negative vorticity. To this end we label our data in the following way. Let l_X, l_Y be the label function on X, Y given by 0, 1 for negative, positive vorticity, respectively. Additionally, due to possible dissipation of vorticity, we focus on unbalanced approaches for the estimation of the transfer operator. We proceed to solve the entropic unbalanced OT problem between μ and ν with respect to the cost function $c(x, y) = d_{\mathrm{E}}(x, y)^2 + d_{\mathrm{E}}^2(l_X(x), l_Y(y))$, regularization parameter $\varepsilon = 0.0003$ and marginal relaxation parameter $\kappa = 0.1$. This can be understood as the

Fig. 8. Left to right: GW transfer kernel, OT transfer kernel. The last two images show distinctly coloured partitions according to the spectral clustering of $K_{\varepsilon,\kappa}^{\mathrm{GW}}$. The plotted numbers are the evaluations of the GW functional at π_{GW} restricted to the respective partitions.

(entropic and unbalanced) Wasserstein distance with an additional penalty on transporting between distinctly signed vorticity. Thus we obtain a solution denoted $\hat{\pi}_{\varepsilon,\kappa}^{\mathrm{W}}$. In the same way, let $\hat{\pi}_{\varepsilon,\kappa}^{\mathrm{GW}}$ be a solution to the unbalanced, fused, entropic GW problem between $(\mathbb{X}, l_X)$ and $(\mathbb{Y}, l_Y)$ and ε, κ as above. On the left-hand side of Figure 8, we illustrate the associated transfer operators $K_{\varepsilon,\kappa}^{\bullet}$, $\bullet \in \{\mathrm{W}, \mathrm{GW}\}$. Similar to the previous examples, the OT transfer operator is not able to recover the underlying rotation. On the other hand, by favoring the preservation of intrinsic distances, the GW transport nicely reflects a counterclockwise rotation. Finally, we apply the discussed spectral clustering procedure, where we focus on $K_{\varepsilon,\kappa}^{\mathrm{GW}}$. To obtain more than two coherent structures, we proceed in a nested manner. More precisely, applying the clustering procedure once yields two partitions of each mm-space X_1, X_2, Y_1, Y_2, respectively. Then we apply the procedure with respect to the associated (labelled) sub mm-spaces

$$\left(\left(X_i, d_X|_{X_i}, \frac{\mu(\cdot \cap X_i)}{\mu(X_i)}\right), l_X|_{X_i}\right), \qquad \left(\left(Y_i, d_Y|_{Y_i}, \frac{\nu(\cdot \cap Y_i)}{\nu(Y_i)}\right), l_Y|_{Y_i}\right),$$

and the restricted transfer operator $K_{\varepsilon,\kappa}^{\mathrm{GW}}|_{X_i \times Y_i}$, $i = 1, 2$. This yields two sub-partitions per partition. We repeat this three times so that we obtain 8 partitions in total. The right-hand side of Figure 8 shows the mm-spaces $\mathbb{X}, \mathbb{Y}$, where points of the same partition are coloured equally. Additionally, we evaluate the GW functional of $\hat{\pi}_{\varepsilon,\kappa}^{\mathrm{GW}}$ restricted to the partitions as explained in Remark 1 and add the evaluation in the plot of $\mathbb{X}$. As expected, the center vortex is clearly identified. Additionally, we are able to identify even smaller structures such as the coherent structures in brown, pink, grey and red. The orange partition attains the smallest GW evaluation by far. This is followed by brown and pink which represent smaller coherent structures and highlights their shape preservation under the transfer $K_{\varepsilon,\kappa}^{\mathrm{GW}}$.

5 Conclusions

In this paper, we proposed a novel approach to estimate dynamical systems based on (unbalanced, fused) GW transport plans. Moreover, we demonstrated that

the obtained transport plans can be leveraged for a spectral clustering procedure to extract coherent structures. The resulting method is convenient as it can be quickly implemented by using out-of-the-box methods for GW and the singular value decomposition. We verified its potency on three numerical examples.

As future work we leave a direct comparison with the method proposed in [2]. The latter provides a numerically more appealing framework for obtaining transport plans which minimize the OT functional under additional invariances such as orthogonal transformations. Moreover, we are interested in applying our method on non-Euclidean data such as e.g. graphs.

References

1. Alvarez-Melis, D., Jaakkola, T, Gromov-Wasserstein alignment of word embedding spaces. In: Proceedings of the 2018 Conference on EMNLP, pp. 1881–1890. ACL, Brussels, Belgium (2018)
2. Alvarez-Melis, D., Jegelka, S., Jaakkola, T., Towards optimal transport with global invariances. In: Proceedings of the 22nd AISTATS, pp. 1870–1879. PMLR, Naha, Okinawa, Japan (2019)
3. Arjovsky, M., Chintala, S., Bottou, L., Wasserstein generative adversarial networks. In: Proceedings of the 34th ICML, pp. 214–223. PMLR, Sydney, NSW, Australia (2017)
4. Beier, F., Beinert, R., Steidl, G., On a linear Gromov–Wasserstein distance. IEEE Trans. Image Process. **31**, 7292–7305 (2022)
5. Beier, F., Beinert, R., Steidl, G., Multi-marginal Gromov–Wasserstein transport and barycenters. arXiv:2205.06725 (2022)
6. Beier, F., von Lindheim, J., Neumayer, S., Steidl, G., Unbalanced multi-marginal optimal transport. J. Math. Imaging Vis. (2022)
7. Bonneel, N., Van De Panne, M., Paris, S., Heidrich, W., Displacement interpolation using Lagrangian mass transport. ACM Trans. Graph. **30**(6), 1–12 (2011)
8. Carlier, G., Ekeland, I., Matching for teams. Econ. Theory **42**(2), 397–418 (2010)
9. Cuturi, M., Sinkhorn distances: Lightspeed computation of optimal transport. In: Advances in Neural Information Processing Systems 26, pp. 2292–2300. Curran Associates, Inc., Lake Tahoe, Nevada, United States (2013)
10. Cuturi, M., Doucet, A., Fast computation of Wasserstein barycenters. In: Proceedings of the 31st ICML, pp. 685–693. PLMR, Bejing, China (2012)
11. Elvander, F., Haasler, I., Jakobsson, A., Karlsson, J., Multi-marginal optimal transport using partial information with applications in robust localization and sensor fusion. Signal Process. **171**, 107474 (2020)
12. Flamary, R., Courty, N., POT Python Optimal Transport library. https://github.com/PythonOT/POT, (2017). Accessed: 19.01.2023
13. Frogner, C., Zhang, C., Mobahi, H., Araya, M., Poggio, T. A., Learning with a Wasserstein Loss. In: Advances in Neural Information Processing Systems 28, pp. 2053–2061. Curran Associates, Inc., Montreal, Canada (2015)

14. Froyland, G., An analytic framework for identifying finite-time coherent structures in time-dependent dynamical systems. Phys. D **250**, 1–19 (2013)

15. Junge, O., Matthes, D., Schmitzer, B.: Entropic transfer operators. arXiv:2204.04901 (2022)

16. Koltai, P., von Lindheim, J., Neumayer, S., Steidl, G.: Transfer operators from optimal transport plans for coherent set detection. Phys. D **426**, 132980 (2021)

17. Kusner, M., Sun, Y., Kolkin, N., Weinberger, K.: From word embeddings to document distances In: Proceedings of the 32nd ICML, pp. 957–966. PMLR, Lille, France (2015)

18. Liero, M., Mielke, A., Savaré, G., Optimal entropy-transport problems and a new Hellinger–Kantorovich distance between positive measures. Invent. Math. **211**(3), 969–1117 (2018)

19. Mémoli, F., Gromov–Wasserstein distances and the metric approach to object matching. Found. Comput. Math. **11**(4), 417–487 (2011)

20. Neumayer, S., Steidl, G.: From optimal transport to discrepancy. In: Chen, K., Schönlieb, CB., Tai, XC., Younces, L. (eds) Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging. Springer, Cham. (2021).

21. Peyré, G., and Cuturi, M., Solomon, J.: Gromov-Wasserstein averaging of kernel and distance matrices. In: Proceedings of the 33rd ICML, pp. 2664–2672. PMLR, New York, NY, United States (2016)

22. Séjourné, T., Feydy, J., Vialard, F.-X., Trouvé, A., Peyré, G. Sinkhorn divergences for unbalanced optimal transport. arXiv:1910.12958 (2019).

23. Séjourné, Th., Vialard, F.–X., Peyré, G., The unbalanced Gromov Wasserstein distance: conic formulation and relaxation. In: Advances in Neural Information Processing Systems 34, pp. 8766–8779. Curran Associates, Inc., Virtual Only (2021)

24. Sturm, K.-T., The space of spaces: curvature bounds and gradient flows on the space of metric measure spaces. arXiv:1208.0434 (2012)

25. Thual, A., Tran, H., Zemskova, T., Courty, N., Flamary, R., Dehaene, S., Thirion, B., Aligning individual brains with fused unbalanced Gromov–Wasserstein. arXiv:2206.09398 (2022)

26. Vayer, T., Chapel, L., Flamary, R., Tavenard, R., Courty, N., Fused Gromov–Wasserstein distance for structured objects. Algorithms **13**(9), 212 (2020)

27. Wang, J., Sesterhenn, J., Müller, W.–C., Coherent structure detection and the inverse cascade mechanism in two-dimensional Navier–Stokes turbulence. arXiv:2203.11336 (2022)

28. Wang, W., Slepčev, D., Basu, S., Ozolek, J., Rohde, G., A linear optimal transportation framework for quantifying and visualizing variations in sets of images. Int. J. Comput. Vis. **101**(2), 254–269 (2013)

29. Xu, H., Luo, D., Zha, H., Carin, L., Gromov–Wasserstein learning for graph matching and node embedding. In: Proceedings of the 36th ICML, pp. 6932–6941 PMLR, Long Beach, California, USA (2019)

30. Zhu, L. and Yang, Y., Haker, S., Tannenbaum, A., An image morphing technique based on optimal mass preserving mapping. IEEE Trans. Image Process. **16**(6), 1481–1495 (2007)