Understanding Mass Transfer Directions via Data-Driven Models with Application to Mobile Phone Data

Alessandro Alla*   Caterina Balzotti†   Maya Briani‡   Emiliano Cristiani‡

Abstract

The aim of this paper is to solve an inverse problem which regards a mass moving in a bounded domain. We assume that the mass moves following an unknown velocity field and that the evolution of the mass density can be described by a partial differential equation, which is also unknown. The input data of the problems are given by some snapshots of the mass distribution at certain times, while the sought output is the velocity field that drives the mass along its displacement. To this aim, we put in place an algorithm based on the combination of two methods: first, we use the Dynamic Mode Decomposition to create a mathematical model describing the mass transfer; second, we use the notion of Wasserstein distance (also known as earth mover's distance) to reconstruct the underlying velocity field that is responsible for the displacement. Finally, we consider a real-life application: the algorithm is employed to study the travel flows of people in large populated areas using, as input data, density profiles (i.e. the spatial distribution) of people in given areas at different time instants. This kind of data are provided by the Italian telecommunication company TIM and are derived by mobile phone usage.

1 Introduction

In this paper we aim at solving an inverse problem which regards a mass moving in a bounded domain with finite velocity. We assume that the mass moves following an unknown velocity field and that the evolution of the mass density can be described by an unknown PDE. The input data of the problem are given by some snapshots of the mass distribution at certain times, while the sought information is the velocity field that drives the mass along its displacement.

The basic idea can be summarized as follows: given two snapshots of the mass distribution at two instants of time, we want to understand where each portion of the mass (which is assumed to be conserved from one instant of time to the other) is transported from/to, i.e. how the first spatial distribution is rearranged in the second one. The goal is pursued by computing a numerical approximation of the Wasserstein distance (also known as earth mover's distance or Mallows distance) between the two consecutive density profiles, specifying a suitable cost function which measures the "energy" consumed by the system for moving forward. The computation of the Wasserstein distance gives, as a by-product, the minimum-cost mass flow from the first to the second configuration, i.e. how the mass distributes in space and time.

Despite the good preliminary results obtained by the above described Wasserstein-based approach [2, 19], the algorithm is found to be excessively expensive both in terms of CPU time and memory

---

*Departamento de Matemática, PUC-Rio, Rio de Janeiro, BR (alla@mat.puc-rio.br).

†Dipartimento di Scienze di Base e Applicate per l'Ingegneria, Sapienza Università di Roma, Rome, IT (caterina.balzotti@sbai.uniroma1.it).

‡Istituto per le Applicazioni del Calcolo "M. Picone", Consiglio Nazionale delle Ricerche, Rome, IT (m.briani@iac.cnr.it, e.cristiani@iac.cnr.it).

requirements. This fact strongly restricts the applicability of the method. To fix this, in this paper we propose to couple the method with the Dynamic Mode Decomposition (DMD): a data-driven technique that takes in input the snapshots of the mass distribution and returns an analytical approximation of the dynamics underlying the mass transfer. More precisely, it provides a system of ODEs which describes the evolution of the mass in any point of the domain. Solving the ODEs, we are able to recover the mass distribution *at any time*, thus increasing at will the number of available snapshots or, analogously, decreasing at will the time frame between them. Controlling the time frame between two consecutive snapshots is the key to simplify the computation of the Wasserstein distance and makes the computation of the flows feasible even on large domains.

Finally, a real-world application of the proposed methodology is illustrated. We are interested in inferring activity-based human mobility flows from mobile phone data. We assume that mobile devices are not singularly tracked, but their logs are aggregated in order to obtain the total number of users in a given area. In this way we get the density profiles (i.e. the spatial distribution) of people in a given area at various instants of time. The dataset we have at our disposal is provided by the Italian telecommunication company TIM. The time frame between two consecutive snapshots is 15 minutes.

As before, the goal is to "assign a direction" to the presence data. In fact, the mere representation of time varying density of people clearly differentiate attractive from repulsive or neutral areas but does not provide any information about the directions of flows of people. In other words, we are interested in a "where-from-where-to" type of information, which reveals travel flows and patterns of people providing a sort of *origin-destination* matrix.

Paper organization. The paper is organized as follows. In Section 2 we present the DMD method and the Wasserstein distance. In Section 3 we show how to couple those methods to obtain an efficient algorithm. In Section 4 we apply the proposed approach to real-life data. In Section 5 we draw our conclusions. Finally, in Appendix A we apply the DMD method to the viscous Burgers' equation.

2 Mathematical background

In this section we recall the building blocks of the methodology proposed in the paper, namely the DMD method used to build a data-driven model, and the Wasserstein distance to determine the transport map driving the moving mass.

2.1 DMD method

DMD is a data-driven method capable of providing accurate assessments of the spatio-temporal coherent structures in a given complex system, or short-time future estimates of such a system. Although a complete list of references for DMD goes beyond the scopes of this work, we would like to mention [3, 5, 6, 13, 15]. In the current work we use the DMD algorithm as in [16].

To begin, we suppose to have a set of data $\mathcal{X} = \{\mathbf{y}(t_0), \ldots, \mathbf{y}(t_{n_{\mathrm{f}}})\}$ for some time instances $\{t_j\}_{j=0}^{n_{\mathrm{f}}}$ with $\mathbf{y}(t_j) \in \mathbb{R}^N, j = 0, \ldots, n_{\mathrm{f}}$ and $\Delta t = t_{j+1} - t_j$ for $j = 0, \ldots, n_{\mathrm{f}} - 1$. The goal of the method is to build a mathematical model upon the dataset $\mathcal{X} \in \mathbb{R}^{N \times (n_{\mathrm{f}}+1)}$. The DMD procedure thus constructs the approximate linear evolution $\widehat{\mathbf{y}}(t)$ for the dataset $\mathcal{X}$ exploiting its low-rank structure:

$$\frac{d\widehat{\mathbf{y}}}{dt} = \widehat{\mathbf{A}}\widehat{\mathbf{y}} \tag{2.1}$$

where $\widehat{\mathbf{A}} \in \mathbb{R}^{N \times N}$ is unknown, $\widehat{\mathbf{y}}(0) = \widehat{\mathbf{y}}_0$, and the solution has the form

$$\widehat{\mathbf{y}}(t) = \sum_{i=1}^{r} \beta_i \boldsymbol{\psi}_i \exp(\omega_i t), \tag{2.2}$$

2

where $r < N$, $\boldsymbol{\psi}_i$ and ω_i are the eigenvectors and eigenvalues of the unknown matrix $\widehat{\mathbf{A}}$. The coefficients β_i of the vector $\boldsymbol{\beta}$ can be determined from the initial data. For example, at $t = t_0$ we have $\mathbf{y}(t_0) = \mathbf{y}_0$ so that (2.2) gives $\boldsymbol{\beta} = \boldsymbol{\Psi}^\dagger \mathbf{y}_0$, where $\boldsymbol{\Psi}$ is a matrix comprised of the DMD modes $\boldsymbol{\psi}_i$ and $\dagger$ denotes the Moore-Penrose pseudo-inverse. To compute the matrix $\widehat{\mathbf{A}}$, we first split the dataset into two snapshot matrices

$$\mathbf{Y} = \begin{bmatrix} | & | & & | \\ \mathbf{y}(t_0) & \mathbf{y}(t_1) & \cdots & \mathbf{y}(t_{n_{\mathrm{f}}-1}) \\ | & | & & | \end{bmatrix}, \quad \mathbf{Y}' = \begin{bmatrix} | & | & & | \\ \mathbf{y}(t_1) & \mathbf{y}(t_2) & \cdots & \mathbf{y}(t_{n_{\mathrm{f}}}) \\ | & | & & | \end{bmatrix} \tag{2.3}$$

and suppose the following linear relation hold true:

$$\mathbf{Y}' = \mathbf{A}\mathbf{Y}, \tag{2.4}$$

where $\mathbf{A} := \exp\left(\widehat{\mathbf{A}}\Delta t\right)$.

Specifically, we assume that $\mathbf{y}(t_j)$ is an initial condition to obtain $\mathbf{y}(t_{j+1})$, i.e. its corresponding output after some prescribed evolution time $\Delta t > 0$. Thus, the DMD method computes the best linear operator $\mathbf{A}$ relating to the matrices above:

$$\mathbf{A} = \mathbf{Y}'\mathbf{Y}^\dagger. \tag{2.5}$$

We refer to $\mathbf{Y}$ and $\mathbf{Y}'$ as input and output snapshot matrices respectively.

The DMD algorithm aims at optimally constructing the matrix $\mathbf{A}$ so that the error between the true and approximate solution is small in a least-square sense, i.e. $\|\mathbf{y}(t) - \widehat{\mathbf{y}}(t)\| \ll 1$. Of course, the optimality of the approximation holds only over the sampling window where $\mathbf{A}$ is constructed, but the approximate solution can be used to make future state predictions, and to decompose the dynamics into various time-scales.

The matrix $\mathbf{A}$ is often highly ill-conditioned and when the state dimension n is large, the aforementioned matrix may be even intractable to analyze directly. Instead, DMD circumvents the eigen-decomposition of $\mathbf{A}$ by considering a low rank representation in terms of a matrix $\widehat{\mathbf{A}}$ projected with the Proper Orthogonal Decomposition (POD). Although the description of the POD method goes beyond the scopes of this paper, we recall that the POD projection solves the following optimization problem

$$\min_{\boldsymbol{\varphi}_1,\ldots,\boldsymbol{\varphi}_r \in \mathbb{R}^n} \sum_{j=0}^{n_{\mathrm{f}}-1} \left\| \mathbf{y}(t_j) - \sum_{i=1}^{r} \langle \mathbf{y}(t_j), \boldsymbol{\varphi}_i \rangle \boldsymbol{\varphi}_i \right\|^2 \quad \text{such that } \langle \boldsymbol{\varphi}_i, \boldsymbol{\varphi}_j \rangle = \delta_{ij}, \tag{2.6}$$

where $\{\boldsymbol{\varphi}_i\}_{i=1}^r$ are the POD projectors. The solution of the optimization problem (2.6) is obtained by means of a Singular Value Decomposition (SVD) of the dataset $\mathbf{Y}$, where the first $r \ll N$ columns of the left singular eigenvectors are the required POD basis. We refer to [18] for a complete description of the POD method. We also mention that POD method is equivalent to Principal Component Analysis (PCA) or Karhunen-Loéve expansion in other contexts (see e.g. [8, 9]).

The exact DMD algorithm proceeds as follows [16]: first, we collect data $\mathbf{Y}, \mathbf{Y}'$ as in (2.3) and compute the reduced, or economy, singular value decomposition of $\mathbf{Y}$

$$\mathbf{Y} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T.$$

We note that the use of the economy SVD is suggested since the matrices $\mathbf{Y}, \mathbf{Y}' \in \mathbb{R}^{n \times n_{\mathrm{f}}}$ with $N \gg n_{\mathrm{f}}$ and that the economy $\mathbf{U} \in \mathbb{R}^{N \times r}$ is sufficient to provide the same approximation of the regular SVD given the limited amount of snapshots.
Furthermore, the DMD typically makes use of low-rank structure so that the total number of modes, $r \ll N$, allows for dimensionality reduction of the dynamical system. Then, we compute the least-squares fit $\mathbf{A}$ that satisfies $\mathbf{Y}' = \mathbf{A}\mathbf{Y}$ and project onto the POD modes $\mathbf{U}$. Specifically, the Moore-Penrose pseudo-inverse of $\mathbf{Y}$ allows us to compute $\mathbf{A} = \mathbf{Y}'\mathbf{Y}^\dagger$, where the Moore-Penrose algorithm provides the least-square fitting procedure. In terms of its low-rank projection, this yields

$$\widehat{\mathbf{A}} = \mathbf{U}^T \mathbf{A} \mathbf{U} = \mathbf{U}^T \mathbf{Y}' \mathbf{V} \boldsymbol{\Sigma}^{-1},$$

and then, we compute the eigen-decomposition of $\widehat{\mathbf{A}} \in \mathbb{R}^{r \times r}$

$$\widehat{\mathbf{A}}\mathbf{W} = \mathbf{W}\mathbf{\Lambda},$$

where $\mathbf{\Lambda}$ are the DMD eigenvalues. Finally, the DMD modes $\mathbf{\Psi}^{\text{DMD}}$ are given by

$$\mathbf{\Psi}^{\text{DMD}} = \mathbf{Y}'\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{W}. \tag{2.7}$$

The algorithm is summarized in Algorithm 1.

Algorithm 1 Exact DMD

Require: Snapshots $\{\mathbf{y}(t_0), \ldots, \mathbf{y}(t_{n_\text{f}})\}$, Time step Δt.
  1: Set $\mathbf{Y} = [\mathbf{y}(t_0), \ldots, \mathbf{y}(t_{n_\text{f}-1})]$ and $\mathbf{Y}' = [\mathbf{y}(t_1), \ldots, \mathbf{y}(t_{n_\text{f}})]$.
  2: Compute the reduced SVD of rank r of $\mathbf{Y}$, $\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$.
  3: Define $\widehat{\mathbf{A}} := \mathbf{U}^T\mathbf{Y}'\mathbf{V}\mathbf{\Sigma}^{-1}$.
  4: Compute eigenvalues and eigenvectors of $\widehat{\mathbf{A}}\mathbf{W} = \mathbf{W}\mathbf{\Lambda}$.
  5: Set $\mathbf{\Psi}^{\text{DMD}} = \mathbf{\Lambda}^{-1}\mathbf{Y}'\mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{W}$.
  6: Set $\omega_i = \frac{\log(\lambda_i)}{\Delta t}$ for (2.2)

In Appendix A we provide a numerical experiment to show the effectiveness of the DMD method. We compute the data from a nonlinear PDE, i.e. the viscous Burgers' equation.

2.2 Wasserstein distance and optimal mass transfer problem

The notion of Wasserstein distance is strictly related to the Monge–Kantorovich optimal mass transfer problem [17], which can be easily explained as follows: given a sandpile with mass distribution ρ^i and a pit with equal volume and mass distribution ρ^f, find a way to minimize the cost of transporting sand into the pit. The cost for moving mass depends on both the distance from the point of origin to the point of arrival and the amount of mass is moved along that path. We are interested in minimizing this cost by finding the optimal path to transport the mass from the initial to the final configuration.

Given two density functions $\rho^\text{i}, \rho^\text{f} : \Omega \to \mathbb{R}$ for some bounded $\Omega \subset \mathbb{R}^n$ such that $\int_{\mathbb{R}^n} \rho^\text{i} = \int_{\mathbb{R}^n} \rho^\text{f}$, we define the L^p-Wasserstein distance between ρ^i and ρ^f as

$$W_p(\rho^\text{i}, \rho^\text{f}) = \left(\min_{T \in \mathcal{T}} \int_\Omega c(\boldsymbol{\xi}, T(\boldsymbol{\xi}))^p \, \rho^\text{i}(\boldsymbol{\xi}) d\boldsymbol{\xi} \right)^{\frac{1}{p}} \tag{2.8}$$

where

$$\mathcal{T} := \left\{ T : \Omega \to \Omega \ : \int_B \rho^\text{f}(\boldsymbol{\xi}) d\boldsymbol{\xi} = \int_{\{\boldsymbol{\xi} : T(\boldsymbol{\xi}) \in B\}} \rho^\text{i}(\boldsymbol{\xi}) d\boldsymbol{\xi}, \quad \forall B \subset \Omega \right\}$$

and $c : \Omega \times \Omega \to \mathbb{R}$ is a given cost function, which defines the cost of transferring a unit mass between any two points in Ω. Note that $\mathcal{T}$ is the set of all possible maps which transfer the mass from one configuration to the other.

It is important to note here that we are not really interested in the actual value of the Wasserstein distance W_p, instead we look for the *optimal map* T^* which realizes the arg min in (2.8), and represents the paths along which the mass is transferred.

2.3 Numerical approximation of the Wasserstein distance

A direct numerical approximation of definition (2.8) is unfeasible, but a discrete approach is still possible. Indeed, we can resort to classical problems (see Hitchcock's paper [7]) and methods (see e.g., [12, Sec. 6.4.1] and [14, Chap. 19]) to recast the original mass transfer problem in the framework of

linear programming (LP). We also refer to [4] for a recent application of this methodology to traffic flow problems.

The idea consists in approximating the set Ω with a structured grid with N square cells $C_1, \ldots, C_N$, as it is commonly done for the numerical approximation of PDEs. We denote by Δx the length of each side of the cells. Then, we define a graph $\mathcal{G}$ whose nodes coincide with the centers of the N cells. Graph's edges are defined in such a way that each node is directly connected with each other, including itself.

Introducing a numerical error (controlled by the parameter Δx), we are allowed to assume that $\forall j = 1, \ldots, N$ all the mass distributed in the cell C_j is concentrated in its center, i.e. in a node of the graph. We come up with an initial mass $m_j^{\mathrm{i}} := \int_{C_j} \rho^{\mathrm{i}} dx$ and a final mass $m_j^{\mathrm{f}} := \int_{C_j} \rho^{\mathrm{f}} dx$, for $j = 1, \ldots, N$, distributed on the graph nodes. Now, we simply aim at optimally rearranging the first mass into the second one moving it among the graph's nodes.

We denote by c_{jk} the cost to transfer a unit mass from node j to node k, and by x_{jk} the (unknown) mass moving from node j to node k. The problem is then formulated as

$$\text{minimize } \mathcal{H} := \sum_{j,k=1}^{N} c_{jk} x_{jk}$$

subject to

$$\sum_k x_{jk} = m_j^{\mathrm{i}} \quad \forall j, \quad \sum_j x_{jk} = m_k^{\mathrm{f}} \quad \forall k \quad \text{and} \quad x_{jk} \geq 0.$$

Defining

$$\mathbf{x} = (x_{11}, x_{12}, \ldots, x_{1N}, x_{21}, \ldots, x_{2N}, \ldots, x_{N1}, \ldots, x_{NN})^{\mathrm{T}},$$
$$\mathbf{c} = (c_{11}, c_{12}, \ldots, c_{1N}, c_{21}, \ldots, c_{2N}, \ldots, c_{N1}, \ldots, c_{NN})^{\mathrm{T}},$$
$$\mathbf{b} = (m_1^{\mathrm{i}}, \ldots, m_N^{\mathrm{i}}, m_1^{\mathrm{f}}, \ldots, m_N^{\mathrm{f}})^{\mathrm{T}},$$

and the $2N \times N^2$ matrix

$$\mathbf{M} = \begin{bmatrix} \mathbb{1}_N & 0 & 0 & \ldots & 0 \\ 0 & \mathbb{1}_N & 0 & \ldots & 0 \\ 0 & 0 & \mathbb{1}_N & \ldots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \ldots & \mathbb{1}_N \\ \mathbf{I}_N & \mathbf{I}_N & \mathbf{I}_N & \mathbf{I}_N & \mathbf{I}_N \end{bmatrix},$$

where $\mathbf{I}_N$ is the $N \times N$ identity matrix and $\mathbb{1}_N = \underbrace{(1 \; 1 \ldots \; 1)}_{N \text{ times}}$, our problem is written as a standard LP problem:

$$\begin{aligned} \min \quad & \mathbf{c}^{\mathrm{T}} \mathbf{x}, \\ \text{subject to} \quad & \mathbf{M} \mathbf{x} = \mathbf{b}, \\ & \mathbf{x} \geq 0. \end{aligned} \tag{2.9}$$

The result of the algorithm is a vector $\mathbf{x}^* := \arg\min \mathbf{c}^{\mathrm{T}} \mathbf{x}$ whose elements x_{jk}^* represent how much mass moves from node j to node k employing the minimum-cost mass rearrangement.

Remark 2.1. The dimension of the LP problem (2.9) is given by the dimensions of the matrix/vectors involved:

$$\dim \mathbf{x} = \dim \mathbf{c} = N^2, \qquad \dim \mathbf{b} = 2N, \qquad \dim \mathbf{M} = 2N^3.$$

Remark 2.2. Hereafter we will refer to problem (2.9) as *global*, in order to stress the fact that it is possible to move the mass from and to any node of the graph.

3 Coupling DMD and optimal mass transfer problem

In this section we describe how we can drastically reduce the size of the LP problem (2.9) by using the DMD method. The resulting algorithm will allow us to study the mass transfer problem on large domains.

3.1 Main idea

The large size of the LP problem (2.9), stressed in Remark 2.1, mainly comes from the fact that the mass is allowed to be transferred from any graph node to any other graph node. However physical constraints prevent this to happening: assuming that the maximal velocity of the mass is $V_{\max}$ and denoting by Δt the time frame from one snapshot to the following one, the maximal distance travelled is $V_{\max}\Delta t$.

In order to reduce the size of the LP problem (2.9), we restrict the set of all possible movements trajectories. The ideal time frame δt of the snapshots would be the one which guarantees that the CFL-like condition (see, e.g., [11, Chapter 14])

$$V_{\max}\delta t \leq \Delta x \tag{3.1}$$

holds true. Indeed, under condition (3.1) the *mass is allowed to move only towards the adjacent cells or not move at all*.

Consider now a generic set of mass distributions $\mathcal{X} = \{\mathbf{m}(t_0), \ldots, \mathbf{m}(t_{n_{\mathrm{f}}})\}$, where $t_j = j\Delta t$, $j = 0, \ldots, n_{\mathrm{f}}$, and Δt is the time frame of the snapshots of the set $\mathcal{X}$. *A priori*, the time step Δt does not necessarily satisfy condition (3.1). In particular, a large distance in time between the snapshots means that there are not enough information for reducing the set of possible movements. However, with the DMD we are able to reconstruct the state of the system at any time instance, even if it is not provided in the original dataset. The coupling of DMD and of optimal mass transfer problem is done in this order:

   i) We first fix $\delta t < \Delta t$ such that it satisfies condition (3.1) and then we reconstruct the solution for each $\widetilde{t}_j = j\delta t$ via DMD, using Algorithm 1. The new set of snapshots we work with is $\widetilde{\mathcal{X}} = \{\mathbf{m}(\widetilde{t}_0), \ldots, \mathbf{m}(\widetilde{t}_{\widetilde{n}_{\mathrm{f}}})\}$, where $\mathbf{m}(\widetilde{t}_j)$ is computed from (2.2), $j = 0, \ldots, \widetilde{n}_{\mathrm{f}}$. We observe that $\widetilde{n}_{\mathrm{f}} > n_{\mathrm{f}}$.

   ii) We recover the flows from the new set $\widetilde{\mathcal{X}}$ by means of an approximation of Wasserstein distance similar to the one done in Section 2.3, but with reduced size, as described in detail below. Note that, since $\widetilde{n}_{\mathrm{f}} > n_{\mathrm{f}}$, we have to solve more LP problems with respect to the global approach, but, despite this, we will get advantages by using this new approach.

Let us denote by d the maximum number of neighbors per cell. The new unknown $\widetilde{x}_{jk}$, corresponding to the mass to be moved from node j to node k, is defined only if j and k are adjacent or j is equal to k. Analogously we define the cost function $\widetilde{c}_{jk}$. We denote by $\widetilde{\mathbf{x}}$ the vector of the unknowns and by $\widetilde{\mathbf{c}}$ the vector associated to the cost function. We introduce the vector $\mathbf{s}$ whose components are indexes of nodes. The first components are the indexes of the nodes adjacent to the node 1, then those adjacent to the node 2 and so until the node N. Vectors $\widetilde{\mathbf{x}}$ and $\widetilde{\mathbf{c}}$ are ordered similarly to $\mathbf{s}$. Since the mass can move only towards a maximum of d nodes, the dimensions of $\widetilde{\mathbf{x}}$ and $\widetilde{\mathbf{c}}$ are lower or

equal than dN. We introduce the matrix

$$\widetilde{\mathbf{M}} = \begin{bmatrix} \mathbb{1}_{r_1} & 0 & 0 & \ldots & 0 \\ 0 & \mathbb{1}_{r_2} & 0 & \ldots & 0 \\ 0 & 0 & \mathbb{1}_{r_3} & \ldots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \ldots & \mathbb{1}_{r_N} \\ & & \widetilde{\mathbf{D}} & & \end{bmatrix},$$

where $\mathbb{1}_{r_i} = \underbrace{(1\ 1 \ldots\ 1)}_{r_i \text{ times}}$, with r_i the number of adjacent nodes for each node i, $r_i \leq d$, and $\widetilde{\mathbf{D}}_{ij} = 1$ if the j-th element of $\mathbf{s}$ is equal to i, otherwise $\widetilde{\mathbf{D}}_{ij} = 0$. Defining the vector $\widetilde{\mathbf{b}} = (m_1^{\mathrm{i}}, \ldots, m_N^{\mathrm{i}}, m_1^{\mathrm{f}}, \ldots, m_N^{\mathrm{f}})^{\mathrm{T}}$ the LP problem becomes

$$\begin{aligned} \min \quad & \widetilde{\mathbf{c}}^{\mathrm{T}} \widetilde{\mathbf{x}}, \\ \text{subject to} \quad & \widetilde{\mathbf{M}} \widetilde{\mathbf{x}} = \widetilde{\mathbf{b}} \\ & \widetilde{\mathbf{x}} \geq 0. \end{aligned} \tag{3.2}$$

Remark 3.1. The dimension of the LP problem (3.2) is given by the dimensions of the matrix/vectors involved:

$$\dim \widetilde{\mathbf{x}} = \dim \widetilde{\mathbf{c}} \leq dN, \qquad \dim \widetilde{\mathbf{b}} = 2N, \qquad \dim \widetilde{\mathbf{M}} \leq 2dN^2.$$

Remark 3.2. Hereafter we will refer to problem (3.2) as *local*, in order to stress the fact that it is possible to move the mass from and to adjacent nodes of the graph only.

3.2 A toy example for the complete algorithm: the advection equation

In this test we propose an example for the complete algorithm described in Section 3.1. Let us consider the advection equation:

$$\begin{cases} \partial_t u(\mathbf{x}, t) + \mathbf{v} \cdot \nabla u(\mathbf{x}, t) = 0 & \mathbf{x} \in \Omega, t \in [0, T] \\ u(\mathbf{x}, t) = 0 & \mathbf{x} \in \partial\Omega \\ u(\mathbf{x}, 0) = u_0(\mathbf{x}) & \mathbf{x} \in \Omega, \end{cases} \tag{3.3}$$

with $\mathbf{x} = (x_1, x_2)$, $\Omega = [-2, 2] \times [-2, 2]$, $u_0(\mathbf{x}) = \max(0.5 - x_1^2 - x_2^2, 0)$ and constant velocity $\mathbf{v} = (0.5, 0.5)$. It is well known that the analytical solution of (3.3) is $u(\mathbf{x}, t) = u_0(\mathbf{x} - \mathbf{v}t)$ provided that we set T small enough to have inactive boundary conditions, as it is the case here. Hereafter we denote by *reference solution* the analytical solution u of (3.3).

Starting from some snapshots of the solution, we aim at reconstructing the velocity field driving the dynamics, i.e. the vector $\mathbf{v} = (0.5, 0.5)$. In doing this, we will also compare the global and the local problem in terms of CPU time, see (2.9) and (3.2) respectively.

We choose a time step δt, and we collect snapshots with a larger temporal step size $\Delta t = \kappa \delta t$ with $\kappa > 1$ and we reconstruct the solution via the DMD method. We note that here the snapshots are computed from the analytical solution of (3.3). In particular, we work on a grid 40×40 and we choose $T = 2$, $\Delta x = 0.1$, $\delta t = 0.05$ and $\Delta t = 2\,\delta t$. We observe that, since $V_{\max} = \|\mathbf{v}\| = \sqrt{2}/2$, this choice of δt fullfills the condition (3.1). The number of snapshots is 40 and the rank r in Algorithm 1 for the DMD reconstruction is 20. Moreover we identify the nodes of the graph $\mathcal{G}$ of the numerical approximation of Wasserstein distance with the cells of the grid.

Figure 1. Reference solution of equation (3.3) at time $t = 0$ on the left and time $t = 2$ on the right.

Choice of the cost function. Since the cost function is used to figure out the "price to pay" for moving the mass from a node of the graph to another one, the most intuitive definition of c is the Euclidean distance in $\mathbb{R}^2$. This choice was proved to be unsuitable for the global problem, see [2]. Indeed, since the global algorithm allows any movements between the nodes of the graph, using the Euclidean distance for the function c we loose the uniqueness of the optimal transfer map T^*. To see this let us assume that we have to move three unit masses one to the right. In the picture on the left of Figure 2 we move the three unit masses of one to the right while in the picture on the right we move only the first mass of three to the right. The Wasserstein distance between the two configuration is clearly the same and equal to three.


Figure 2. Three small movements versus one large movement.

The solution proposed in [2] to fix this issue was to force the minimization algorithm to select primarily the small movements penalizing the large ones. To get this, the cost function was defined as

$$c(\boldsymbol{\xi}_1, \boldsymbol{\xi}_2) = \|\boldsymbol{\xi}_1 - \boldsymbol{\xi}_2\|_{\mathbb{R}^2}^{1+\varepsilon}, \tag{3.4}$$

where $\boldsymbol{\xi}_1$ and $\boldsymbol{\xi}_2$ are the coordinates of the nodes of the graph and $\varepsilon > 0$ is a small parameter which accounts for the penalization.

In Figure 3 we show the level sets of the solution to (3.3) together with some arrows indicating the reconstructed velocity field $\mathbf{v}$. More precisely, on the left column of figure we show the results obtained with the local algorithm (3.2), on the center with the global algorithm with Euclidean distance and on the right with the global algorithm with the cost function defined in (3.4) (with $\varepsilon = 0.1$). Similarly, the panels on the top show the results obtained with the reference solution whereas the panels on the bottom those obtained with the DMD solution.

The local algorithm and the global one with the correction in (3.4) are able to reconstruct the velocity field $\mathbf{v} = (0.5, 0.5)$ with accurate approximation. The global algorithm with Euclidean distance, instead, is less precise, since the optimal flow does not correspond to the velocity field. Moreover,
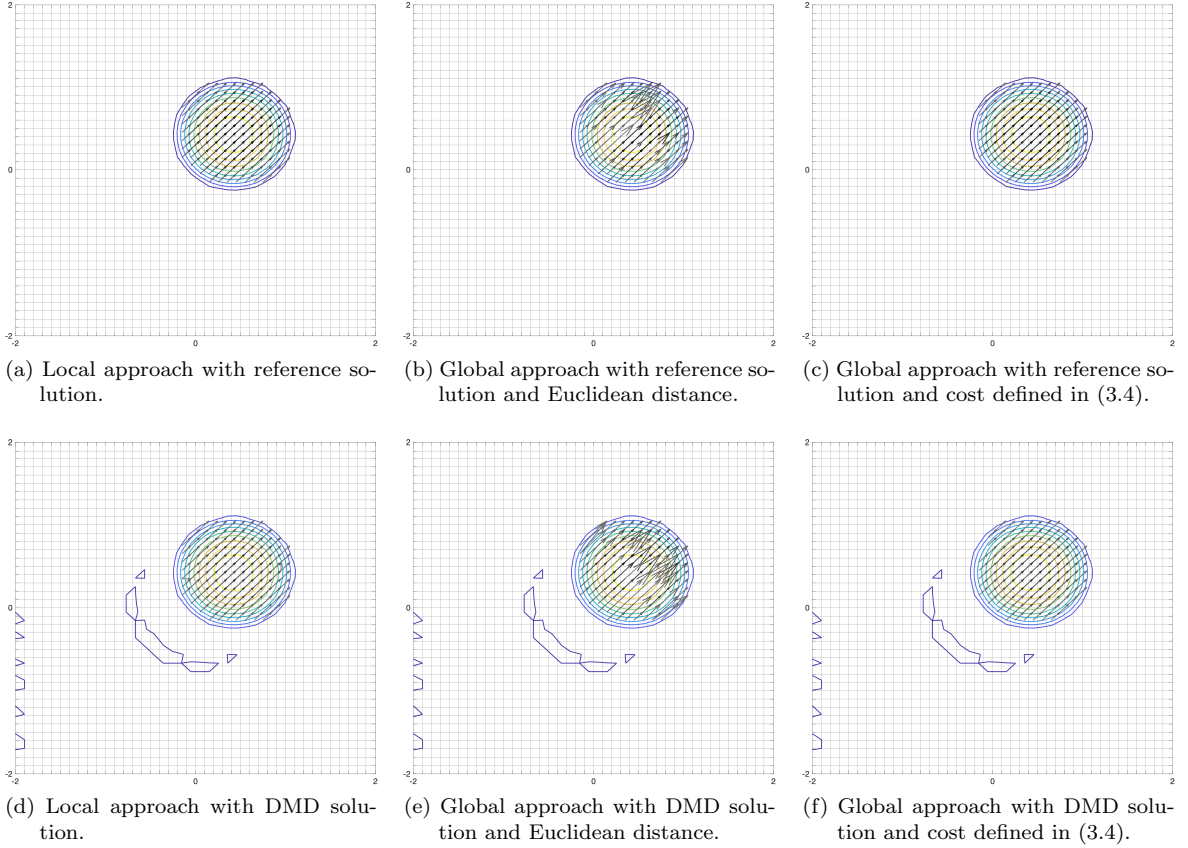
(a) Local approach with reference solution.

(b) Global approach with reference solution and Euclidean distance.

(c) Global approach with reference solution and cost defined in (3.4).

(d) Local approach with DMD solution.

(e) Global approach with DMD solution and Euclidean distance.

(f) Global approach with DMD solution and cost defined in (3.4).

Figure 3. Reconstructed flows between $t_1 = 0.725$ and $t_2 = t_1 + \delta t$ superimposed to the level sets of the solution to (3.3) at t_2. Left column: local algorithm. Central column: global algorithm with c equal to the Euclidean distance. Right column: global algorithm with c defined in (3.4). Top row: the reference solution to (3.3) is used for computation. Bottom row: the DMD reconstruction of the solution to (3.3) is used for computation.

the algorithm based on the DMD reconstruction of the solution introduces small oscillations in the solution to (3.3). This is expected in such hyperbolic problems since the decay of the singular values of the dataset is very slow and the initial condition is non-smooth. However, such oscillations do not have an effect in the reconstruction of the flow.

To further validate our approach we compute the numerical error of the proposed algorithm. Since the results obtained with the global algorithm (2.9) with the Euclidean distance as cost function are the least accurate, we focus only on the other two approaches. For each time step $t_n = n\delta t$ we define $\mathbf{x}_E^*(t_n)$ and $\widetilde{\mathbf{x}}_E^*(t_n)$ as the solutions to the LP problems (2.9) and (3.2) respectively at time t_n, when the known terms $\mathbf{b}(t_n)$ and $\widetilde{\mathbf{b}}(t_n)$ are chosen as the reference solution to (3.3) at time t_n. Analogously, $\mathbf{x}_D^*$ and $\widetilde{\mathbf{x}}_D^*$ are the solution when the known terms are obtained with the DMD method. The vectors $\mathbf{x}_E^*(t_n)$, $\widetilde{\mathbf{x}}_E^*(t^n)$, $\mathbf{x}_D^*(t^n)$ and $\widetilde{\mathbf{x}}_D^*(t_n)$, for $n = 1, \ldots, \lfloor \frac{T}{\delta t} \rfloor + 1$, are finally collected as the columns of the matrices $\mathbf{X}_E^*$, $\widetilde{\mathbf{X}}_E^*$, $\mathbf{X}_D^*$ and $\widetilde{\mathbf{X}}_D^*$ respectively. We define the errors

$$E^G := \frac{\|\mathbf{X}_E^* - \mathbf{X}_D^*\|_F}{\|\mathbf{X}_E^*\|_F}, \qquad E^L := \frac{\left\|\widetilde{\mathbf{X}}_E^* - \widetilde{\mathbf{X}}_D^*\right\|_F}{\left\|\widetilde{\mathbf{X}}_E^*\right\|_F}, \tag{3.5}$$

9

where $\|\cdot\|_F$ is the Frobenius norm. In Table 1 we compare the errors defined in (3.5) obtained from the simulations on a grid with $N \times N$ nodes, for $N = 20$, 30 and 40. As we can see from the table, increasing the number of nodes we reduce the space step $\Delta x = 4/N$ and the time step $\Delta t = \Delta x/2$ and we obtain the decrease of the error between the reference solution and the DMD reconstruction. Moreover, the error obtained with the local algorithm is significantly smaller than the one obtained with the global approach.

Table 1. Comparison of the errors defined in (3.5).

N	Δx	Δt	E^G	E^L
20	0.20	0.100	0.185	0.030
30	0.13	0.067	0.159	0.028
40	0.10	0.050	0.122	0.020

Finally, in Table 2 we compare the computational time between the global approach (2.9), with the cost function as in (3.4), and the local approach (3.2) with respect to the nodes of the graph. We observe that the local algorithm is always faster than the global one. The difference between the two approaches becomes more relevant when we refine the grid by increasing the number of nodes, and thus the number of time steps. Specifically, for a grid 40×40, the local algorithm required a few seconds whereas the global one more than three hours.

Table 2. Computational time.

N	Global Exact	Global DMD	Local Exact	Local DMD
20	18.10 s	18.89 s	0.28 s	0.36 s
30	11 min	11 min	0.93 s	1.41 s
40	3 h 6 min	3 h 7 min	2.40 s	4.85 s

4   Application to real mobile phone data

In this section we focus on a specific application of the proposed approach. The real dataset gives information about the spatial distribution of people in a large populated area. The goal is to understand the travel flows of people, focusing in particular on recurring patterns and daily flows of commuters.

4.1   Dataset

The Italian telecommunication company TIM provides estimates of mobile phones presence in a given area in raster form: the area under analysis is split into a number of elementary territory units (ETUs) of the same size (about 130×140 m^2 in urban areas). The estimation algorithm does not singularly recognize users and does not track them using GPS. It simply counts the number of phone attached to network nodes and, knowing the location and radio coverage of the nodes, estimates the number of TIM users within each ETU at any time. TIM has now a market share of 30% with about 29.4 million mobile lines in Italy (AGCOM, Osservatorio sulle comunicazioni 2/2017).

The data we considered refer to the area of the province of Milan (Italy), which is divided in 198,779 ETUs, distributed in a rectangular grid 389×511. Data span six months (February, March and April 2016 and 2017). The entire period is divided into time intervals of 15 minutes, therefore we have 96 data per day per ETU in total. In Figure 4 we graphically represent presence data at a fixed time. We observe that the peak of presence is located in correspondence of Milan city area. Figure 5 shows the presences in the province of Milan in a typical working day in the left panel. The

Figure 4. 3D-plot of the number of TIM users in each ETU of Milan's province on April 18, 2017.

curve in the image decreases during the night, it increases in the day-time and decreases again in the evening. These variations are due to two main reasons: first, the arrival to and departure from Milan's province of visitors and commuters. Second, the fact that when a mobile phone is switched off or is not communicating for more than six hours, its localization is lost. The presence value that most represents the population of the province is observed around 9 pm., when an equilibrium between traveling and phone usage is reached. This value changes between working days and weekends, but it is always in the order of 1.3×10^6. On the right panel of Figure 5 we show the trend of presence data


Figure 5. Trend of presences in the province of Milan during a typical working day (left). Trend of presences in the province of Milan during April 2017 (right).

during April 2017. We can observe a cyclical behavior: in the working days the number of presences in the province is significantly higher than during the weekends. It is interesting to note the presence of two low-density periods on April 15-18 and on April 22-26, 2017, determined respectively by the Easter and the long weekend for the Italy's Liberation Day holiday.

4.2 DMD approach on TIM data

As explained in Section 2.1, we can reconstruct the presence data in each cell at any time. We denote by $\mathbf{m}(t^{\mathrm{i}}) \in \mathbb{R}^N$ the vector containing the number of people present in the N cells at a certain quarter of an hour and by $\mathbf{m}(t^{\mathrm{i}}+15\,\mathrm{min}) \in \mathbb{R}^N$ the same quantities at the consecutive time step. Applying the DMD method we are able to calculate $\mathbf{m}(t) \in \mathbb{R}^N$ for any $t \in [t^{\mathrm{i}}, t^{\mathrm{i}}+15\,\mathrm{min}]$, see (2.2). We also note that DMD can be applied to higher dimensional dataset through randomized methods as explained in [1].

To validate the DMD approach on the TIM dataset we study the daily error using only half of the data at our disposal in the DMD reconstruction. More precisely, we denote by $\mathbf{P}^{\mathrm{data}}$ the matrix whose columns correspond to the real data of presences stored every 15 minutes. Then, we reconstruct the data of presences every minute with the DMD algorithm, using snapshots every 30 minutes. In other words, we exploit only one column out of two of $\mathbf{P}^{\mathrm{data}}$ to build the matrix $\widetilde{\mathbf{P}}^{DMD}$, which collects the reconstructed data every minute. Since a day contains 96 quarters of an hour and 1440 minutes, $\mathbf{P}^{\mathrm{data}}$ is a matrix $96 \times N$, while $\widetilde{\mathbf{P}}^{DMD}$ is a matrix $1440 \times N$. Finally, to compare the real data with the DMD reconstruction, we extract from $\widetilde{\mathbf{P}}^{DMD}$ the rows corresponding to the original interval of 15 minutes into the matrix $\widehat{\mathbf{P}}^{\mathrm{DMD}}$, of dimensions $96 \times N$, and then we define the error as:

$$E = \frac{\left\| \mathbf{P}^{\mathrm{data}} - \widehat{\mathbf{P}}^{\mathrm{DMD}} \right\|_F}{\left\| \mathbf{P}^{\mathrm{data}} \right\|_F}. \tag{4.1}$$

In Figure 6 we show the daily error (4.1) for an entire month of data in the whole area of the province of Milan. The daily error is of order 10^{-2}, which certifies the accuracy of the DMD method.


Figure 6. Daily error defined in (4.1) where DMD has been computed with $r = 95$ in step 2 of Algorithm 1.

4.3 Understanding human mobility flows

Following the approach described in Section 3, we define the graph $\mathcal{G}$ by exploiting the subdivision of the area of the province of Milan (Italy) into ETUs; we identify the N nodes of the graph with the corresponding center of the ETUs, ordered from the left to the right and from the top to the bottom. The result is a rectangular graph $\mathcal{G}$ divided into N_R rows and N_C columns ($N = N_R \times N_C$). The mass $m_j(t_n)$ is defined as the average number of presences in the node j at time $t_n = n\,15\,\mathrm{min}$.

Let us assume that $V_{\max}$ in (3.1) is equal to 50 km/h. Since the dimensions of the ETUs is around 150 m, to apply the DMD we fix the new time step δt equal to 10 seconds. With this choice we assume

that people can move only towards the eight adjacent nodes of the rectangular graph, or not move at all. We observe that the mass in the nodes on the four corners of the graph can move only towards four directions (adjacent nodes or no movement), while the mass along the boundaries can move only towards six directions. In this way, the total number of possible movements between the cells is given by

$$\widetilde{N} = \underbrace{4 \cdot 4}_{\text{corners}} + \underbrace{6 \cdot 2 \left(N_R + N_C - 4\right)}_{\text{boundaries}} + \underbrace{9 \left(N - 4 - 2 \left(N_R + N_C - 4\right)\right)}_{\text{internal nodes}} < 9N. \tag{4.2}$$

The vectors $\widetilde{\mathbf{x}}$ and $\widetilde{\mathbf{c}}$ associated to the unknown moving mass and to the cost function have length $\widetilde{N}$, while the matrix $\widetilde{\mathbf{M}}$ of the LP problem (3.2) has dimension $2N \times \widetilde{N}$. In Table 3, we compare the dimensions of the vectors and the matrices of the two different LP problems: it is clear that the computational time to solve problem (3.2) is significantly reduced respect to problem (2.9).

Table 3. Comparison between dimensions of matrices and vectors for the two methods.

Algoritm	Vectors Dimension	Matrix Dimension
Global	N^2	$2N^3$
Local	$\widetilde{N} \, (< 9N)$	$2N \times \widetilde{N} \, (< 18N^2)$

Choice of the cost function c. Since the ETUs are rectangles of length $\ell_x = 130$ m and $\ell_y = 140$ m, we define the cost function c_{jk} for the local algorithm as follows:

$$c_{jk} = \begin{cases} 0 & \text{if } j = k \text{ or } j \text{ and } k \text{ are not adjacent} \\ \ell_x & \text{if } j \text{ and } k \text{ are horizontally adjacent} \\ \ell_y & \text{if } j \text{ and } k \text{ are vertically adjacent} \\ \sqrt{\ell_x^2 + \ell_y^2} & \text{if } j \text{ and } k \text{ are diagonally adjacent.} \end{cases}$$

For the global approach, we use the cost function defined in (3.4) with $\varepsilon = 0.1$.

To sum up, in order to solve the mass transfer problem for a whole day using snapshots taken every 15 minutes (real data) we have to solve 96 global LP-problems (2.9), whereas with the DMD algorithm we have to solve 8640 local LP-problems (3.2). As we will see in the following section, despite the larger number of LP problems, the local approach is more convenient than the global one.

Remark 4.1. The Wasserstein distance is defined between two distributions of equal mass (see (2.8)). In our case the conservation of mass between two consecutive snapshots is not guaranteed. Let us consider a couple of snapshots with different total mass $\sum_j m_j^0 \neq \sum_j m_j^1$. To correctly apply the algorithms for the identification of flows, we compute the mass in excess between the two snapshots and then we uniformly distribute it in all the nodes of the graph with lower mass. A more sophisticated approach could be the one suggested in [10], where a definition of Wasserstein distance between two distributions with different mass is proposed.

4.4 Numerical results

In this section we show the results obtained with the local algorithm (3.2) to study the flows of commuters and the influence of great events on human mobility. In both cases the rank r in step 2 of Algorithm 1, used for the construction of the DMD solution, is 95. The flows are represented by arrows; we draw only those which correspond to the most significant movements, and we associate a darker color to the arrows corresponding to a larger movement of people. For graphical purposes, in the following plots we aggregate 6 time steps δt to show 1-minute mass transport.

4.4.1 Flows of commuters

In the following simulations we consider the area of the Province of Milan during a generic working day. Milan is one of the biggest Italian city and it attracts many workers from outside. The city of Milan is located in the right part of the Province, therefore we mainly see movements from/to the left part of the analyzed area. In the top panel of Figure 7 we show the morning flows of a working day: we clearly see that the arrows are directed towards the city of Milan. In particular, we zoom on the arrows which overlap the roads heading to Milan. In the bottom panel of Figure 7 we show the opposite phenomenon: in the evening people go away from work to come back home inside the Province of Milan.

CPU times. Considering data for a 6 hours frame on an area of 144×240 ETUs the local algorithm requires 144 hours of CPU time and works with 360 snapshots. The global approach (2.9) is not able to analyze such an area, since the matrix $\mathbf{M}$ in (2.9) has a computationally unmanageable dimension.

4.4.2 Flows influenced by a great event

In this test we show how the algorithm is able to capture the way a big event influences human flows. The event we have considered is the exhibition of the *Salone del Mobile*, held every April at Fiera Milano exhibition center in Rho, near Milan. We analyze a square area of 31×31 ETUs centered around Fiera Milano. In the left panel of Figure 8 we show the morning fluxes directed to the exhibition area whereas in the right panel we show the evening flows directed from the exhibition area to the outside.

CPU times. For a simulation of 18 hours, from 06:00 to 23:45, on an area of 31×31 ETUs, the local algorithm requires 6 minutes of CPU time while the global approach requires 30 minutes. Furthermore, we observe that the local algorithm works with 1065 snapshots of data, whereas the global one with 266 snapshots.

5 Conclusions

In this paper we have proposed an efficient method to solve an inverse mass transfer problem, consisting in recovering the dynamics underlying the mass displacement. The proposed algorithm prevents to handle large displacements of the mass, thus saving CPU time and memory allocation with respect to other recently proposed solutions.

The application of the methodology to a real dataset describing the movement of people was also investigated. It is useful to note here that the applicability of the Wasserstein-based methodology was not at all obvious, since it is based on assumptions which are not totally fullfilled. Indeed, the choice of the cost function c does not take into account the fact that people move mainly along roads, are stopped by obstacles, buildings, etc., and in general are not free to move in all directions. Moreover, and most important, the computation of the Wasserstein distance stems from a global optimization problem in which the mass is considered as a whole. In other words, the optimal transport map T^* is found by minimizing the cost of the displacement of the whole crowd, without any distinction among people, and with no regards about individual targets. Despite this, the results we have obtained (see especially those in Figure 7 and 8) are exactly as one can expect, meaning that the method is, overall, robust enough to work well even if the constitutive assumptions are not totally fullfilled.
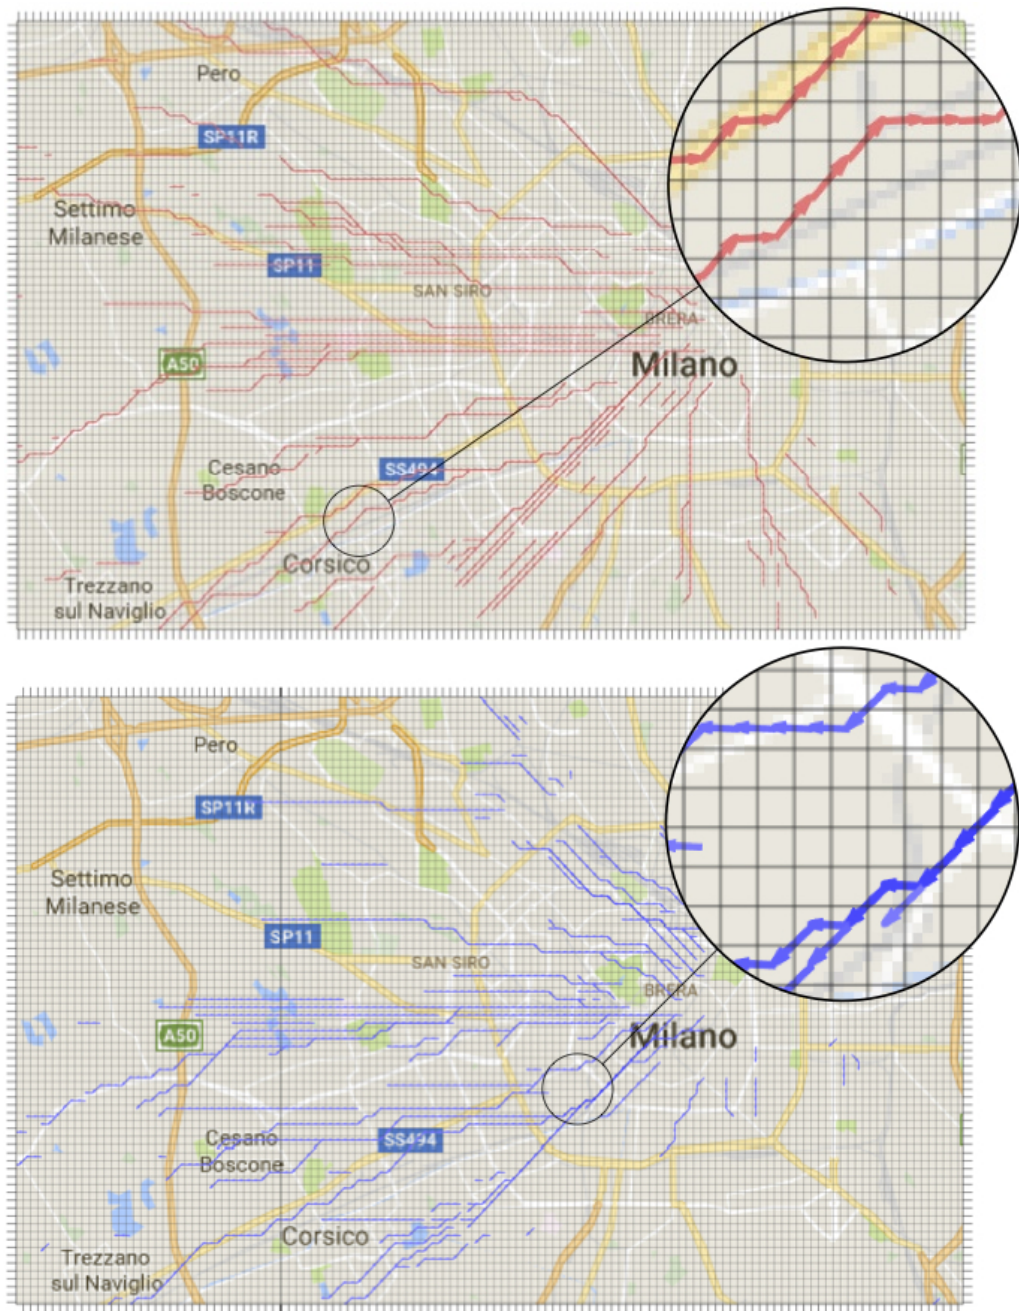
Figure 7. Flows of commuters during the morning (09:00-09:01) of a generic working day (top). Flows of commuters during the evening (18:00-18:01) of a generic working day (bottom).
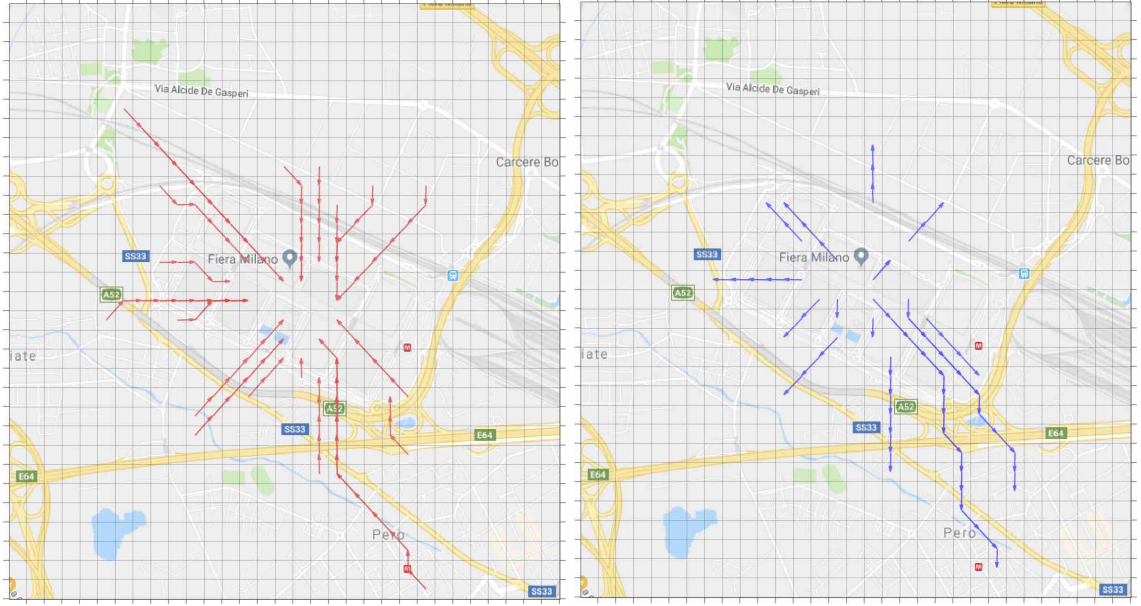
Figure 8. Main flows directed to the exhibition area in the morning (09:45 - 09:46) (left), Main flows from the exhibition area in the evening (18:00 - 18:01).(right)

A   The DMD for the viscous Burgers' equation

Here we propose a numerical experiment using data generated by the 2D viscous Burgers' equation:

$$\begin{cases} \partial_t y - \varepsilon \Delta y + y(\partial_{x_1} y + \partial_{x_2} y) = 0 & \mathbf{x} \in \Omega, t \in [0, T] \\ y(\mathbf{x}, t) = 0 & \mathbf{x} \in \partial\Omega \\ y(\mathbf{x}, 0) = \sin(\pi x_1) \sin(\pi x_2) & \mathbf{x} \in \Omega, \end{cases} \tag{A.1}$$

with $\mathbf{x} = (x_1, x_2)$, $\Omega = [0, 1] \times [0, 1]$, $T = 1$ and $\varepsilon = 0.01$. The numerical approximation of equation (A.1) has been carried out by a finite difference method on a square grid with $N = 1600$ nodes, using an adaptive Runge-Kutta method for the time integration with temporal step size $\Delta t > 0$. We refer to [11] for further details on the numerical method. In what follows, since we do not know the analytical solution of equation, we consider as reference solution of (A.1) its numerical approximation with $\Delta t = 0.0125$.

To show how DMD reconstructs the model we consider as dataset $\mathcal{X}$ the numerical approximation of (A.1) for the following temporal step sizes $\Delta t = \{0.1, 0.05, 0.025, 0.0125\}$ with snapshots matrices of dimension $N \times 11, N \times 21, N \times 41, N \times 81$, respectively. Once the snapshots have been computed we apply Algorithm 1 to reconstruct the solution corresponding to a fixed $\delta t = 0.0125$. Here, we aim at testing the capability of the DMD method to approximate the dataset for the missing information.

In Figure 9, we show the reference solution of (A.1) in the top row and the absolute difference between the reference solution and the reconstructed one with DMD in the bottom row for two time instances. We note that the dataset for the DMD reconstruction corresponds to the solution of (A.1) with $\Delta t = 0.1$.

Then, in Figure 10 we show the relative error between the DMD approximation and the reference solution of (A.1) as follows:

$$E^{DMD} := \frac{\left\| \mathbf{Y}^R - \mathbf{Y}^{DMD} \right\|_F}{\left\| \mathbf{Y}^R \right\|_F}, \qquad \widehat{E}^{DMD}(t) := \frac{\left\| \mathbf{y}^R(t) - \mathbf{y}^{DMD}(t) \right\|_2}{\left\| \mathbf{y}^R(t) \right\|_2}, \tag{A.2}$$
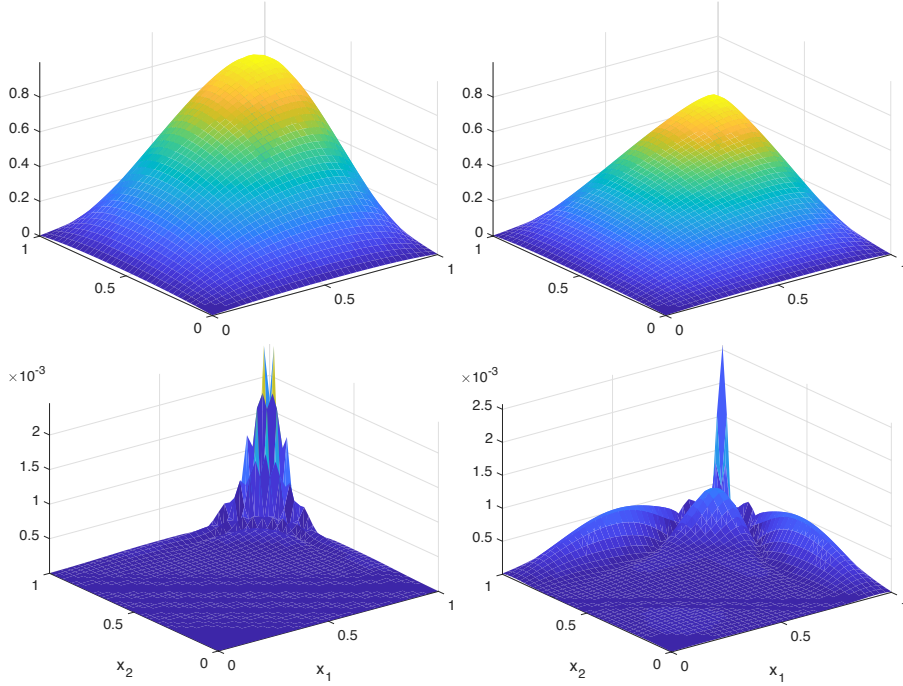
Figure 9. Top: Reference Solution for Burgers' equation at time $t = 0.4875$ (left) and $t = 0.9875$ (right). Bottom: absolute difference between reference solution and DMD reconstruction with data $\Delta t = 0.1$ at time $t = 0.4875$ (left) and $t = 0.9875$ (right).

where $\mathbf{y}^R(t)$ is the reference solution and $\mathbf{y}^{DMD}(t)$ is its DMD approximation. The notation $\mathbf{Y}^R, \mathbf{Y}^{DMD}$ refers to the matrices where each column contains the reference solution and its DMD approximation, respectively, for different time instances.


Figure 10. Error analysis of the DMD reconstruction with a reference solution computed with $\Delta t = 0.0125$ and dataset corresponding to $\Delta t = \{0.1, 0.05, 0.025, 0.0125\}$ as in the $x-$axis using E^{DMD} in (A.2) (left), relative error $\widehat{E}^{DMD}(t)$ in (A.2) for each time instance for the DMD reconstruction with $\Delta t = \{0.1, 0.05, 0.025\}$ (middle) and singular values of the snapshots matrix with $\Delta t = 0.0125$ together with the error behavior E^{DMD} for different rank (right).

It is clear that the more information on the system, e.g. amount of snapshots, the better error behavior for the DMD solution (see left panel in Figure 10). However, we can already obtain accurate solutions with a poor dataset. In the middle panel of Figure 10 we show the error of the model for each time instance for the dataset corresponding to $\Delta t = \{0.1, 0.05, 0.025\}$ using $\widehat{E}^{DMD}(t)$ defined in (A.2). We note that DMD behaves very well for time instances provided in the snapshot set. This

is to show that DMD is able to reconstruct accurately the solution for missing information. In this example the rank r in the DMD method is 9. Finally, we would like to emphasize that, although the DMD model is linear, it works even for nonlinear problem as, e.g., (A.1). For the sake of completeness we also show in the right panel of Figure 10 the decay of error E^{DMD} in (A.2) modifying the rank r in the DMD approximation. As expected the error decays as the rank increases. The singular values of snapshots set are also in the picture.

References

[1] A. Alla and J.N. Kutz. Randomized model order reduction. *Advances in Computational Mathematics*, https://doi.org/10.1007/s10444-018-09655-9, 2019.

[2] C. Balzotti, A. Bragagnini, M. Briani, and E. Cristiani. Understanding Human Mobility Flows from Aggregated Mobile Phone Data. In *15th IFAC Symposium on Control in Transportation Systems CTS 2018*, volume 51-9, pages 25–30, 2018.

[3] D.A. Bistrian and I.M. Navon. An improved algorithm for the shallow water equations model reduction: Dynamic Mode Decomposition vs POD. *Internat. J. Numer. Methods Fluids*, 78:552 – 580, 2015.

[4] M. Briani, E. Cristiani, and E. Iacomini. Sensitivity analysis of the LWR model for traffic forecast on large networks using Wasserstein distance. *Commun. Math. Sci.*, 16(1):123–144, 2018.

[5] P. Héas and C. Herzet. Optimal kernel-based Dynamic Mode Decomposition. https://arxiv.org/pdf/1710.10919.pdf.

[6] P. Héas and C. Herzet. Optimal low-rank Dynamic Mode Decomposition. https://arxiv.org/pdf/1701.01064.pdf.

[7] F. L. Hitchcock. The distribution of a product from several sources to numerous localities. *Stud. Appl. Math.*, 20(1-4):224–230, 1941.

[8] I.T. Jolliffe. *Principal component analysis*. Springer, 2002.

[9] K. Pearson. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2:559 – 572, 1901.

[10] Benedetto Piccoli and Francesco Rossi. Generalized Wasserstein distance and its application to transport equations with source. *Arch. Ration. Mech. Anal.*, 211(1):335–358, 2014.

[11] A. Quarteroni and A. Valli. *Numerical Approximation of Partial Differential Equations*. Springer, 1994.

[12] F. Santambrogio. *Optimal transport for applied mathematicians*, volume 87 of *Progress in Nonlinear Differential Equations and their Applications*. Birkhäuser/Springer, Cham, 2015. Calculus of variations, PDEs, and modeling.

[13] P. J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, 2010.

[14] S.M. Sinha. *Mathematical Programming: Theory and Methods*. Elsevier, 2005.

[15] G. Tissot, L. Cordier, N. Benard, and B.R. Noack. Model reduction using dynamic mode decomposition. *Comptes Rendus Mécanique*, 342:410–416, 2014.

[16] J.H. Tu, C.W. Rowley, D.M. Luchtenburg, S.L. Brunton, and J.N. Kutz. On dynamic mode decomposition: Theory and applications. *AIMS*, 1:391–321, 2014.

[17] C. Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.

[18] S. Volkwein. *Model Reduction using Proper Orthogonal Decomposition*. Lecture notes, University of Konstanz, 2011.

[19] D. Zhu, Z. Huang, L. Shi, L. Wu, and Y. Liu. Inferring spatial interaction patterns from sequential snapshots of spatial distributions. *International Journal of Geographical Information Science*, 32(4):783–805, 2018.