Coarse-grained particle model for pedestrian flow using diffusion maps

Christian Marschler[*] and Jens Starke[†]
Department of Applied Mathematics and Computer Science, Technical University of Denmark, Matematiktorvet 303B, DK-2800 Kongens Lyngby, Denmark

Ping Liu[‡] and Ioannis G. Kevrekidis[§]
Department for Chemical and Biological Engineering, Princeton University, Princeton, New Jersey 08544, USA
(Received 7 August 2013; published 14 January 2014)

Interacting particle systems constitute the dynamic model of choice in a variety of application areas. A prominent example is pedestrian dynamics, where good design of escape routes for large buildings and public areas can improve evacuation in emergency situations, avoiding exit blocking and the ensuing panic. Here we employ diffusion maps to study the coarse-grained dynamics of two pedestrian crowds trying to pass through a door from opposite sides. These macroscopic variables and the associated smooth embeddings lead to a better description and a clearer understanding of the nature of the transition to oscillatory dynamics. We also compare the results to those obtained through intuitively chosen macroscopic variables.

DOI: 10.1103/PhysRevE.89.013304 PACS number(s): 07.05.Kf, 89.20.−a, 02.70.−c

I. INTRODUCTION

Understanding the dynamics of interacting particle models is desirable in many contexts, including the modeling of pedestrian crowds, a subject of high social relevance. This type of knowledge can be used for the design of emergency exits and evacuation strategies, and to improve the flow of large crowds of pedestrians in nonemergency situations, e.g., at airport security, large conferences, or shopping malls. Pedestrians can be usefully described as particles, interacting with each other via so-called social forces as well as with the environment (see Refs. [1,2] and [3–6] for further references on pedestrian and many-particle models). This description naturally leads to the study of systems of ordinary differential equations with a large number of dimensions, i.e., four per particle (two positions and two velocities on the plane). In Ref. [7], a typical scenario where two pedestrian crowds try to pass through a door from opposite sides [1] is studied: Using the door width as a bifurcation parameter, a critical door width can be found, at which the system undergoes an apparent macroscopic Hopf bifurcation: the system transitions from a blocked state to an oscillating state with pedestrians from the two sides alternating in crossing through the door. One can simulate this system at the level of interacting individuals, yet the dynamics exhibit an inherent separation of time scales, suggesting the possibility to successfully describe the system by coarse-graining methods. The derivation of a successful collective motion model in terms of good macroscopic variables has, however, proven to be difficult [7].

In general, many dynamical systems (including interacting particles, like our pedestrian model) are characterized by a separation of scales. In such cases, it is sometimes possible to describe and understand the motion of large ensembles of particles in a collective sense through the derivation of explicit equations for the relevant leading statistics (e.g., moments). At intermediate system sizes, the infinite particle limit assumptions that underpin such collective equations become inaccurate, yet separation of time scales and the associated convergence of the high-dimensional dynamics to a low-dimensional, slow manifold (cf. Fenichel's theory [8], see also the slaving principle [9,10]) are still present. It then becomes crucial to determine good sets of macroscopic observables that parametrize this manifold; these are the variables in terms of which collective equations can now be formulated (closed). The same issue arises in equation-free methods [11,12], where a so-called restriction operator maps the high-dimensional variables to a useful low-dimensional representation. Macroscopic observables that are good candidates for low-dimensional descriptions may be known from experience, e.g., low-order moments of a particle distribution, yet for many new systems it is desirable to find such variables in a systematic and algorithmic fashion based on simulation or experimental data.

This links our particle modeling with machine learning (and, in particular, with manifold learning techniques): given observations of trajectories of our model we must determine how many, and then which slow variables are sufficient to describe the collective dynamics of pedestrian flow.

Classical linear approaches for gaining insight into large-scale, complicated data structures include principal component analysis (PCA) [13] and multidimensional scaling (MDS) [14,15]. These methods have been successfully applied to numerous problems in physics and chemistry (see, e.g., the monograph [16] and references therein). PCA and MDS, being linear methods by construction, since they employ linear combinations of the data set, are not able to economically recover complicated nonlinear structures in a data set. This has led to the invention and development of nonlinear data-mining techniques such as isomap [17], local linear embedding (LLE) [18], and spectral methods [19,20].

[*]chrms@dtu.dk
[†]jsta@dtu.dk
[‡]Present address: Department of Molecular, Cellular, and Developmental Biology, Yale University, New Haven, Connecticut 06520, USA; ping.liu@yale.edu
[§]yannis@princeton.edu; also Program in Applied and Computational Mathematics.

TABLE I. Model parameters.

Number of pedestrians N_p	200
Terminal velocity v^0	1.5 m s^{-1}
Reaction time τ	0.22 s
Pedestrian-pedestrian repulsion V	15 m^2 s^{-2}
Pedestrian-wall repulsion V_B	10 m^2 s^{-2}
Pedestrian-pedestrian length scale σ	1 m
Pedestrian-wall length scale R	2 m
Corridor width C_w	5 m
Corridor length C_l	45 m

In this paper, the recently developed nonlinear manifold learning method of diffusion maps [21,22] is used to analyze pedestrian model data, focusing on the oscillatory regime, at a macroscopic level. Systematically selecting a few leading diffusion map components naturally leads to user-independent, data-based dimension reduction; this is particularly helpful in problems for which experience and empirical insight in appropriate variable selection is lacking.

The remainder of the paper is structured as follows. The pedestrian model and the diffusion map algorithm are introduced in Secs. II and III, respectively. In Sec. IV the diffusion map representation of the pedestrian model is computed and the results are compared with those in Ref. [7]. Section V contains a brief discussion and an outlook for future applications.

II. PEDESTRIAN MODEL

The particle model from Ref. [7] is used to study the behavior of pedestrians in a long, narrow corridor, trying to pass through a door, modeled as an opening in the middle of the corridor, from opposite sides. The observed behavior drastically depends on the door width w. A small door blocks the corridor, while a larger door gives rise to oscillatory behavior, with crossings from alternate sides, in an apparent Hopf bifurcation. The model is a variation of the model of Helbing and Molnár [1,2], who studied the behavior of pedestrians affected by social forces, extended by noise in order to avoid deadlock situations.

According to the social force model, the equations of motion for a particle (a pedestrian) is given as

$$\ddot{\mathbf{z}}_i = \mathbf{F}_i^0 + \sum_j \mathbf{f}_{ij} + \sum_B \mathbf{f}_{iB} + \mathbf{n}_i, \tag{1}$$

where $\mathbf{z}_i \in \mathbb{R}^2$ is the ith particle position $i = \{1, \ldots, N_p\}$, and N_p is the total number of particles. The four different contributions in (1) are the direction force $\mathbf{F}_i^0$, the pedestrian-pedestrian interaction $\mathbf{f}_{ij}$, the pedestrian-wall interaction $\mathbf{f}_{iB}$, and the noise $\mathbf{n}_i$. They are described in more detail in Appendix A. All model parameters can be found in Table I. Throughout the paper, numerical values are given in units of the characteristic length scale for the pedestrian interaction σ [(A2)].

Since pedestrians interact with the environment via the wall force $\mathbf{f}_{iB}$, the geometry of the corridor plays an important role. In Ref. [7] a corridor of length C_l and width C_w with a door of width w in the middle of the corridor is studied (Fig. 1). Two

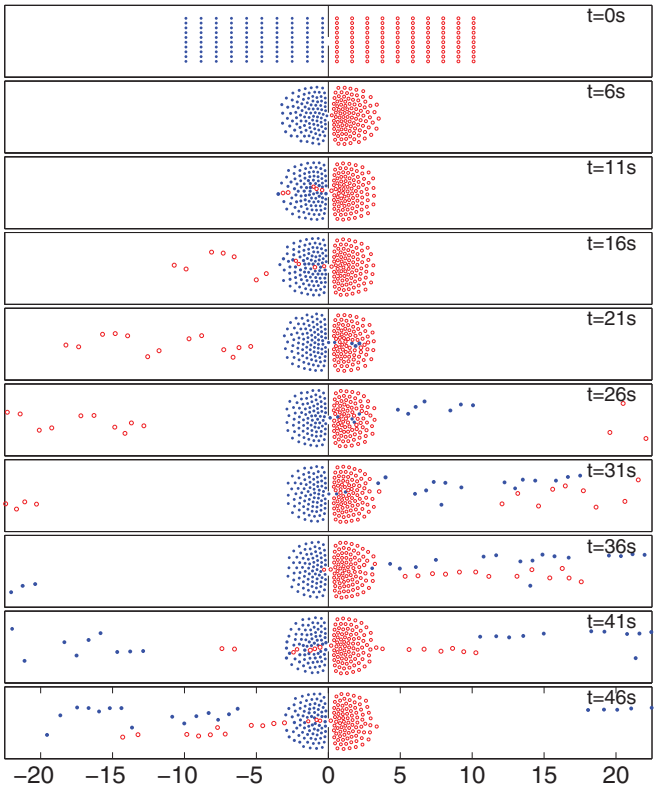

FIG. 1. (Color online) Snapshots showing the oscillatory dynamics of the pedestrian motion for $w = 0.6 > w^* = 0.55$. After initialization, the pedestrians form two milling crowds one on each side of the door ($t = 6$ s). At $t = 11$ s, the (red circle) crowd on the right begins moving through the door, leading to a net flux to the left. After some time, pressure from the right-hand side of the door has decreased enough so that the (blue dot) crowd on the left can break through ($t = 21$ s). At $t = 36$ s, the situation is reversed, and now the red circle crowd can start moving through the door again. This behavior repeats in a periodic fashion, analyzed in Ref. [7].

crowds of pedestrians try to pass the door from opposite sides and the two crowds have the same population size $N_p/2$. For pedestrians starting on one side of the door, the target direction $\mathbf{e}_i^0$ is the center of the door. At the moment pedestrian i passes the door, $\mathbf{e}_i^0$ is updated to point to the end of the corridor, i.e., $\mathbf{e}_i^0$ is chosen parallel to the corridor longitudinal axis. The number of pedestrians is conserved by applying periodic boundary conditions.

In Ref. [7], a coarse equation-free bifurcation analysis was presented using the macroscopic variables

$$m = \frac{1}{2}(m_{\text{red}} + m_{\text{blue}}) \quad \text{and} \quad \dot{m} = \frac{dm}{dt}, \tag{2}$$

where

$$m_\alpha = \frac{\sum_{i\in\alpha} \kappa(x_i) x_i}{\sum_{i\in\alpha} \kappa(x_i)} \tag{3}$$

is a weighted average for pedestrian group α (red, blue) using the weighting function $\kappa(x)$, which gives more weight to pedestrians close to the door. The particular choice of

macroscopic variables was a compromise between filtering for noise reduction and clarity of the ensuing macroscopic description. The more natural choice $m = (m_{\text{red}} - m_{\text{blue}})/2$ and $\dot{m}$ as macroscopic variables gave rise to a noisy macroscopic signal. Further analysis by tools from dynamical systems theory, e.g., bifurcation analysis, was hindered by this noise. Therefore, (2) has been chosen to study the pedestrian behavior (see Ref. [7] for a detailed discussion). An apparent Hopf bifurcation is detected at a door width w^*. For $w < w^*$, the door is too small for pedestrians to pass. They gather in front of the door in what (neglecting small fluctuations of $\mathbf{n}_i$) corresponds to a macroscopic equilibrium. For $w > w^*$, the door is large enough to let pedestrians pass through, leading to macroscopic oscillations (Fig. 1). Moreover, the fast time scale is observed during the early stages of the transient after initialization ($t = 0$). These transients decay over a time scale of $\approx$3 s, while the oscillations have a longer time scale of 40 s. For systems with a smaller gap between time scales, transients take longer. Those systems can still be analyzed by means of diffusion maps using longer simulations.

The selection of variables in Ref. [7] was intuitive and may well depend on the authors. The purpose of the following is to find macroscopic variables intrinsic to the problem data; we also expect these variables to result in good noise filtering.

III. DIMENSION REDUCTION BY DIFFUSION MAPS

Diffusion maps were recently proposed (see Refs. [21–23]) as a nonlinear manifold learning/dimension reduction technique. The goal is to find a (nonlinear) coordinate transformation (i.e., a diffusion map), between the data space and a (low-dimensional) embedding space; the Euclidean distance in the embedding space approximates the *diffusion distance* (defined in Appendix B). Briefly, and qualitatively, the diffusion distance between two data points is small if it is easy to transition between them in a well-defined diffusion process on a graph determined by the data (see Appendix B for details).

If the high-dimensional data happen to lie on (close to) a low-dimensional curved manifold, diffusion maps have been used to extract a parametrization of this manifold and gain insight into the geometric structure underlying the data (see, e.g., Refs. [24–30]). Details of the procedure for diffusion map computation, highlighting the dimension reduction aspect, are presented in Appendix B; they are intended mainly to introduce notation, and are not an extensive review on diffusion maps.

We focus on data resulting from dynamic simulations/observations (possibly multiple ones) of dynamical systems (here, our particle-based pedestrian model); the diffusion maps computation is a purely postprocessing step. We ignore the temporal structure of the data (which also allows us to merge different trajectories); for the exploitation of time information in the form of delay reconstructions in a diffusion map context see Ref. [31]. The ordering in time of data points $x_i = x(t_i)$, taken from a trajectory $x(t)$ at discrete times $t_i = i\,\Delta t, i \in \{1, \ldots, N\}$, will thus not influence the diffusion map construction. The only quantities of interest are the pairwise distances of the points in data space; it is thus possible to analyze the structure of data when the underlying equations of motion are unknown.

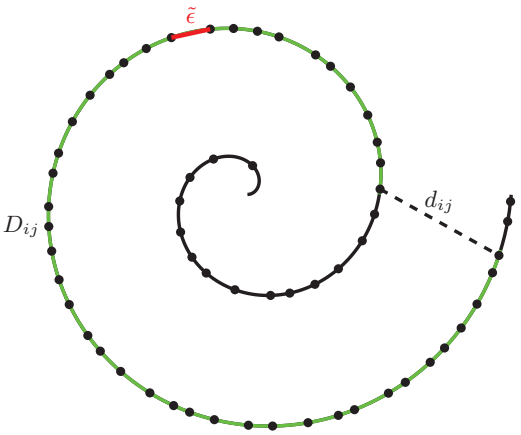

FIG. 2. (Color online) One-dimensional manifold embedded in $\mathbb{R}^2$. The data points are created using (4). $\tilde{\epsilon}$ is representative of a characteristic distance between data points on the manifold [red (dark gray) segment]. d_{ij} denotes the Euclidean distance between two points i and j (dashed line). Although the Euclidean distance is small between these points, the geodesic distance along the manifold, denoted by D_{ij}, is much larger [green (light gray) segment], making it a much better measure of the actual closeness of the two points.

A standard toy illustration example is the swiss roll data set in Fig. 2 (cf. [32]). An embedding into $\mathbb{R}^2$ is given by

$$(x_1, x_2) = (\theta \cos\theta, \theta \sin\theta), \tag{4}$$

where $\theta \in [0, 4\pi]$. Assume that the two-dimensional data points on this sampled manifold (black dots in Fig. 2) are the result of experimental observations or of a dynamical simulation. Although this manifold lies in two dimensions, it is only one-dimensional, i.e., both coordinates are functions of a single variable, θ, that parameterizes the curve. Clearly, any linear dimension reduction method, e.g., PCA and MDS, would fail to detect the one-dimensional structure of the manifold: projection of the data in Fig. 2 on any line would mix the order of data points on the manifold. Techniques such as diffusion maps, as we briefly outline below, can successfully perform this reduction.

To determine the intrinsic data geometry from such a data set X, diffusion maps use Markov chains to describe a diffusion process *on the data set*. Pairwise Euclidean distances d_{ij} are computed for all data points, and weighted (soft-thresholded, using the scaling parameter ϵ) through the diffusion kernel A_{ij} [(B3) in the Appendix B] to give pairwise affinities between the points. The rows of the resulting affinity matrix A are normalized to yield a Markov transition matrix M between points in the data set [(B4)]. The time-t diffusion distance D_t is defined in terms of this Markov matrix; it is small, if the t-step transition probability in the Markov chain is high [(B5)]. The transformation [(B10)] ensures that the Euclidean distance in diffusion map space is (an approximation of) the diffusion distance between the data points [(B7)].

For the procedure to yield informative results several (often problem-dependent) considerations apply; if, for example, in Fig. 2 the scaling parameter ϵ in (B3) is not comparable to some characteristic distance $\tilde{\epsilon}$ between the data, but even larger than the d_{ij} shown, points will be identified as close neighbors that

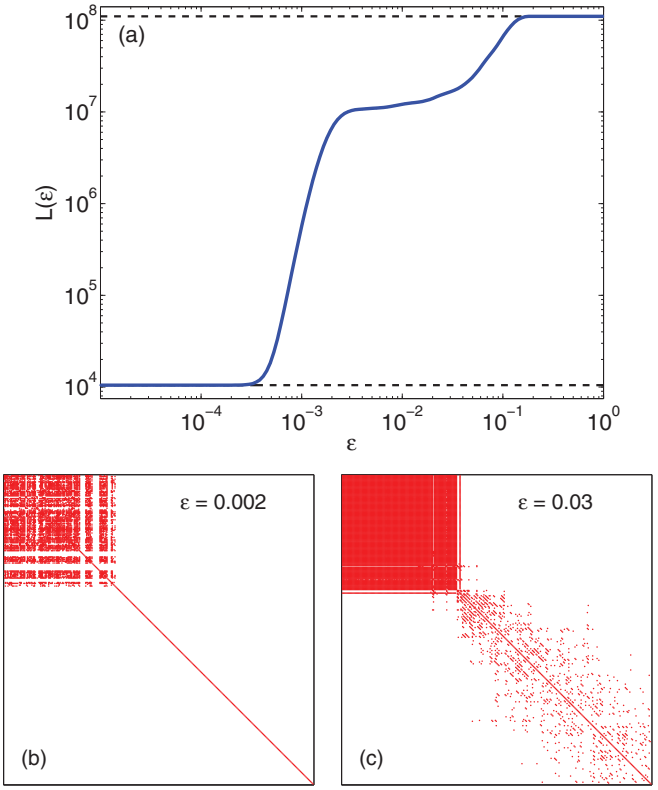

FIG. 3. (Color online) Determining appropriate ϵ values for observations over a range of door widths ($w \in [0.5,0.7]$ with $\Delta w = 0.01$ and 500 observations for each w). (a) The number L of elements in the distance matrix d smaller than ϵ is plotted as a function of ϵ. The limiting behavior (5) is shown as the dashed black lines. In between these limits, $L(\epsilon)$ shows two scaling regimes, namely $[4 \times 10^{-4}, 2 \times 10^{-3}]$ and $[2 \times 10^{-2}, 2 \times 10^{-1}]$. To clarify this, the sparsity pattern of the distance matrix d is inspected for two values of ϵ, $\epsilon = 2 \times 10^{-3}$ representing the blocked regime ($w < w^*$) and $\epsilon = 0.03$ representing the oscillating regime ($w > w^*$). (b) and (c) show the parameter-ordered matrix elements (small w: top left, large w: bottom right), which are smaller than ϵ (dots). For $\epsilon = 0.002$ [(b)], the dots are all located in the upper left corner, denoting blocked states. The oscillating regime is visible at $\epsilon = 0.03$ [(c)]. Oscillating state distances do not register in the left panel, since they are larger than $\epsilon = 0.002$.

intuitively should not be. Selection of an appropriate value for ϵ in the kernel (B3) is also problem/data dependent; one can arrive at such a value through investigation of the scaling behavior of the pairwise distances d_{ij}. Let $L(\epsilon)$ be the number of d_{ij} smaller than ϵ. For large ϵ all d_{ij} are smaller, while for small enough ϵ only the diagonal d_{ii} is smaller (actually, zero):

$$\lim_{\epsilon\to\infty} L(\epsilon) = N^2, \quad \lim_{\epsilon\to 0} L(\epsilon) = N; \tag{5}$$

clearly, ϵ should be chosen between those limits. This behavior is illustrated for the pedestrian flow example in Fig. 3, and will be discussed in more detail below. From scaling arguments, the dimension of the manifold can be estimated from the slope of $L(\epsilon)$ on a logarithmic plot [33].

Another important consideration is the relative scaling of the data coordinates so that the (weighted) Euclidean distance in the numerator of the diffusion kernel is informative for the problem considered; this would arise for example in chemical composition data where different components are present at proportions differing by orders of magnitude.

Diffusion map computation for large data sets can be computationally expensive. Naive, nonsparse storage of the Markov matrix grows like $\mathcal{O}(N^2)$. The computation of the $k+1$ largest eigenvectors can be performed taking advantage of the sparsity of the matrix d (cf. [31]). When the diffusion map coordinates for a new data point $\mathbf{x}_{N+1}$ are needed, adding the point to the data set and repeating all diffusion map computations from scratch is not necessary. Indeed, one can estimate the diffusion map coordinates of this $N+1$st out-of-sample point using the Nyström extension [25,34]: Using (B10), the $N+1$st coordinate of eigenvector $\boldsymbol{\Psi}_i$ is approximated as

$$\boldsymbol{\Psi}_{i,N+1} = \frac{1}{\lambda_i} \sum_{l=1}^{N} M_{N+1,l} \boldsymbol{\Psi}_{i,l}, \tag{6}$$

for $i = 1,\ldots,k$. It is only necessary to compute a single new row of M, thus saving computation time.

IV. RESULTS

We now turn to the application of the diffusion map algorithm (Sec. III) to specific scenarios of pedestrian flow (Sec. II). The section is structured as follows: First a metric on the high-dimensional data space is introduced, in order to construct the distance matrix in (B2). Then we use diffusion maps to obtain a two-dimensional embedding of the dynamics for fixed door width, $w = 0.7 > w^*$, and compare to the results in Ref. [7]. Finally, a diffusion map embedding for data points assembled for different values of w is computed. This allows us to characterize the bifurcation and compare the Hopf bifurcation points obtained from the two approaches.

A. Pedestrian model data

A data point $\mathbf{x}_j$, taken from a trajectory of (1), contains the positions and velocities of all pedestrians

$$\mathbf{x}_j = [\mathbf{Z}_1,\ldots,\mathbf{Z}_{N_p}] \in \mathbb{R}^n, \tag{7}$$

where $\mathbf{Z}_k = [z_{x,k}, z_{y,k}, \dot{z}_{x,k}, \dot{z}_{y,k}]$ for pedestrian k, and $n = 4N_p$. In a preprocessing step, the data points are transformed componentwise to the interval $[0,1]$ to make them comparable. More precisely, let $G = \{x, y, v_x, v_y\}$ denote the set of positions and velocities and $C = \{\text{red}, \text{blue}\}$ denote the color (grayscale), i.e., crowd, of the pedestrian, respectively. We choose the pedestrian labels such that $I_{\text{red}} = \{1,\ldots,N_p/2\}$ and $I_{\text{blue}} = \{N_p/2+1,\ldots,N_p\}$ denote the two index sets for the red (circles) and blue (dots) pedestrians, respectively. In a first transformation step, we shift the minimum to zero by

$$\tilde{z}_{g,i} = z_{g,i} - \min_{g,c} \quad \forall i \in I_c, \quad g \in G, \quad c \in C, \tag{8}$$

where $\min_{g,c} = \min_{i\in I_c} z_{g,i}$ for $g \in G, c \in C$. Afterwards, the data is scaled with

$$\tilde{\tilde{z}}_{g,i} = \frac{\tilde{z}_{g,i}}{\max_{g,c}} \quad \forall i \in I_c, \quad g \in G, \quad c \in C, \tag{9}$$

where $\max_{g,c} = \max_{i \in I_c} \tilde{z}_{g,i} \neq 0$ for $g \in G, c \in C$, onto the interval $[0,1]$. We drop the tilde in the following for notational convenience.

In the new data set X, the pedestrians are still labeled. Our selection of a pairwise distance between configurations should be invariant to particle label permutations. One could use, for example, an earth-mover's distance [35] for this purpose; we choose instead a metric easier to compute, employing the mean and centered moments of four features of the data set: red circle/blue dot x, y positions and red circle/blue dot velocities. A new data vector

$$\tilde{\mathbf{x}}_i = [\langle Z \rangle_{g,c}, m_2(Z)_{g,c}, \ldots, m_{100}(Z)_{g,c}] \quad \forall g \in G, \quad c \in C, \tag{10}$$

is constructed, where $\langle \cdot \rangle_{g,c}$ and $m_i(Z)_{g,c}$ is shorthand notation for the mean and ith centered moment in a list of all feature-color combinations, respectively. The resulting data set is defined as in (B1)

$$X = \{\tilde{\mathbf{x}}_i \in \mathbb{R}^n | i = 1, \ldots, N\}, \tag{11}$$

where $n = 800$. Note, that the dimension of the data points is conserved. This pre-processing step is *only* applied to compute the distance between snapshots without labels of the pedestrians. The distance between two observations $\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j \in X$ can then be defined as

$$d_{ij} = \|\tilde{\mathbf{x}}_i - \tilde{\mathbf{x}}_j\|. \tag{12}$$

Diffusion maps will allow us to reduce this 800-dimensional data set to a two-dimensional representation of the dynamics.

B. Length scale selection

Before the algorithm from Sec. III can be applied, an estimate for the length scale ϵ is required. The function $L(\epsilon)$ [(5)] is investigated for data points obtained over a range of w values (Fig. 3). In this figure, the limiting behavior for $\epsilon \to \infty$ and $\epsilon \to 0$ is clearly visible. Between the two limiting plateaus, a third one exists for $\epsilon \in [2 \times 10^{-3}, 3 \times 10^{-2}]$. It separates two scaling regimes. The scaling regime for smaller ϵ is representative of blocked state dynamics, when the particles simply jiggle on each side of the door, and their configurations do not change appreciably. The larger ϵ regime is representative of the oscillatory dynamics, where appreciable variations between successive data snapshots arise. This is echoed in the size of the elements d_{ij} plotted in matrix format (bottom panel in Fig. 3). In order to rationalize the different regimes, the matrix elements of d are ordered by the door width parameter, i.e., small w observations are found to the top left and large w observations towards the bottom right of Fig. 3 (bottom panel), respectively. The small distances [smaller than ϵ, marked in red (dark gray)] between blocked state configurations originate from model noise. Note that distances between blocked and oscillating states are typically large [empty regions at the top right and bottom left corner in Fig. 3 (bottom panel), respectively]. After inspection of this figure, we chose the fixed value $\epsilon = 0.03$ for all subsequent computations.

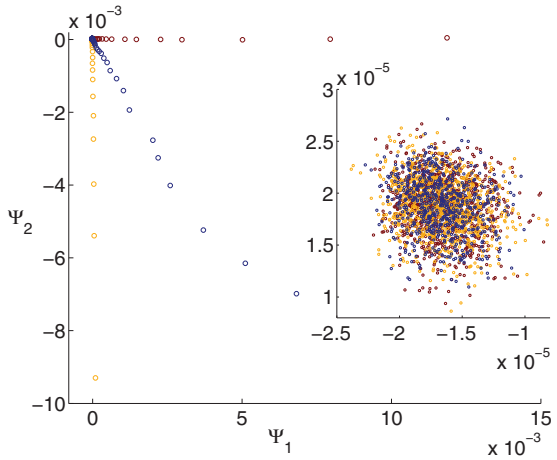

FIG. 4. (Color online) diffusion map embeddings for three trajectories in the blocked regime ($w = 0.4 < w^*$), including transients. The inset magnifies the region of the (noisy) final state. Color (grayscale) indicates the different trajectories of the system.

C. Diffusion map embedding

We first analyze, using diffusion maps [and, in particular, the one-step diffusion, i.e., $t = 1$ in (B5) and (B10) from Appendix B], the blocked regime for small door widths $w < w^*$. Here, the pedestrians congregate on both sides of the door without being able to pass through it. Initial transients in diffusion map coordinates for this regime are shown in Fig. 4 for 100 s and $\Delta t = 0.1$ s. Figure 4 shows the two-dimensional embedding from the blocked regime ($w < w^*$). The axis labels follow the naming in previous literature, where Ψ_j on the axes denotes the jth component of $\hat{\mathbf{y}}$ [(B10)], i.e., the component in the direction corresponding to the jth largest eigenvalue (after the trivial one at 1). It is clear in Fig. 4 that, after an initial transient, the data points become randomly distributed in diffusion map space; the spread of the ball in the inset is indicative of the noise in the simulation.

The picture changes drastically when analyzing a trajectory for, say, $w = 0.7 > w^*$, all other simulation parameters being kept constant. Now, a particle simulation over 500 s with $\Delta t = 0.1$ s is used, resulting in 5000 observations. Initial transients are ignored in the remainder by using the last 3500 data points only. The resulting two- and three-dimensional diffusion map embeddings are shown in Fig. 5.

The third component can be written as a function $\Psi_3(\Psi_1, \Psi_2)$, and therefore a two-dimensional embedding is sufficient to describe the long-term dynamics. These results will later be compared to the original ones in Ref. [7] in Sec. IV D.

Finally, the diffusion map embedding for data points taken from trajectories over a range of door widths $w \in [0.5, 0.7]$ (sampled with $\Delta w = 0.01$) is computed. Each trajectory was now computed for 5000 time steps and subsampled: every third data point of the last 1500 time steps was used for the embedding. Since the periodic orbits have a period of about 40 s, this covers approximately four periods and therefore yields enough data for a good embedding. The resulting matrix size (10500×10500) is close to the limit of what can be routinely handled by MATLAB on a workstation [36] without using

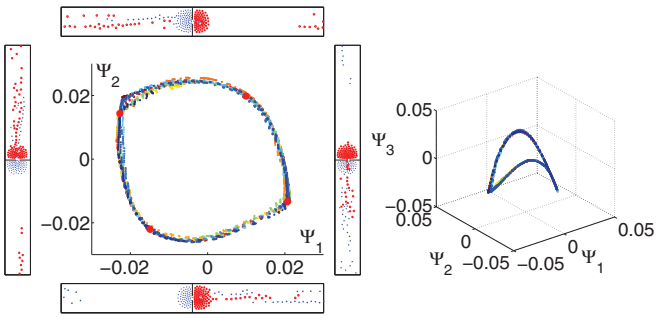

FIG. 5. (Color online) diffusion map embedding of data in the oscillating regime ($w = 0.7 > w^*$). The last 3500 data points of a trajectory (sampled with $\Delta t = 0.1$) are shown. A periodic orbit is clearly visible in the $\Psi_1-\Psi_2$ plane. Big red (dark gray) dots denote points in diffusion map space for which characteristic microscopic states are included as surrounding insets in the figure. Ψ_3 in the three-dimensional embedding can be written as a function of Ψ_1 and Ψ_2.

special algorithms for memory optimization. The resulting embeddings are shown in Fig. 6. All possible two-dimensional projections as well as the full three-dimensional embedding for the first three eigenvectors are shown.

The color (grayscale) here encodes the corresponding door width of the data points [blue (black and light gray): $w = 0.5$, red (dark gray): $w = 0.7$]. It is clearly visible in the (Ψ_2,Ψ_3) plane, that a transition from a blocked state [blue (black) points at (0,0)] to an oscillating state occurs. The amplitude of the oscillations increases with door width, in accordance with

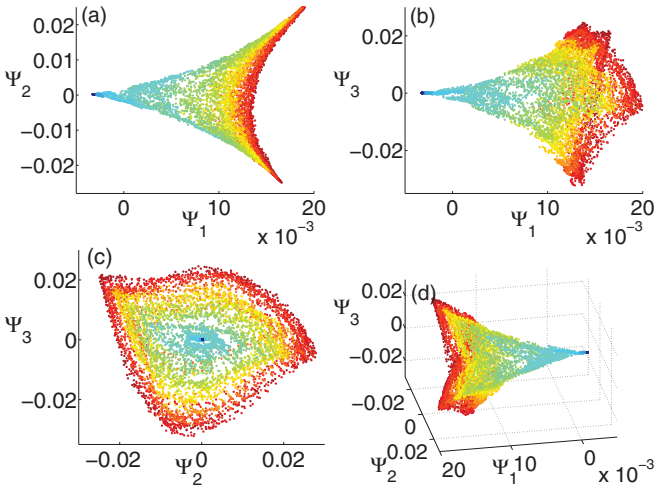

FIG. 6. (Color online) Diffusion map embeddings computed from a data set containing points for different parameter values w. The parameter range for the door width is 0.5 to 0.7 with $\Delta w = 0.01$. Every third data point from the last 1500 iterations for each parameter value has been used, resulting in a data set with 10500 data points. Various projections of the data set are shown in (a)–(d) in diffusion map space. The color (grayscale) encodes the door width from $w = 0.5$ [dark blue (black and light gray)] to $w = 0.7$ [red (dark gray)]. The Ψ_1 coordinate roughly encodes the w regime. The increase in oscillation amplitude is best visible in the (Ψ_2,Ψ_3) projection [(c)], showing periodic solutions for several door widths. As in Ref. [7], the amplitude of the oscillations grows with increasing w. The full three-dimensional representation is shown in (d).

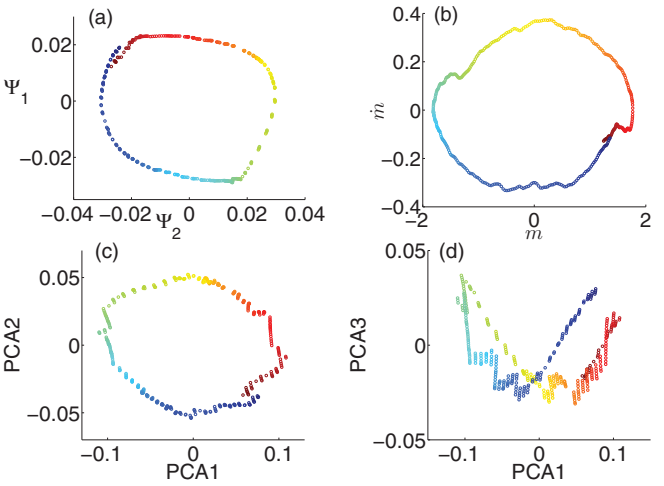

FIG. 7. (Color online) Coarse phase portraits using (a) diffusion maps, (b) the intuitive variables of Ref. [7] and (c), (d) principal components. Data points along a limit cycle are shown for the last 400 time steps of a 5000 time step simulation. Color (grayscale) indicates time (see text). There is a clear one-to-one correspondence between $(m,\dot{m})$ and (Ψ_1,Ψ_2). The diffusion map embedding appears to remove the corners in (b) [red (dark gray) and green (light gray) dots]. Two-dimensional projections of phase portraits in terms of leading PCA coefficients (c) and (d) are also included for comparison. The two-dimensional diffusion map embedding has the smoothest appearance.

the findings in Ref. [7]. One can argue, from inspection of the different projections, that Ψ_1 encodes the parameter w, distinguishing small door widths (small values of Ψ_1) from large ones (large values of Ψ_1).

D. Comparison with other macroscopic representations

We now compare the macroscopic representation of Ref. [7], which they denoted $(m,\dot{m})$, with the diffusion map representation (Ψ_1,Ψ_2). The last 400 time steps of a 5000 time step simulation, covering one oscillation period for a door width $w = 0.7$ are used for this comparison in Fig. 7. Diffusion maps do not consider time labeling of the data points, yet to assist the interpretation of these results (and since, here, the data come from a single long trajectory, so that the time labels are actually available) we use color (grayscale) in these embeddings to encode time. Blue (black) data points lie at the beginning of the trajectory, and the color (grayscale) progresses over green (light gray) to red (dark gray).

There is a clear one-to-one correspondence between $(m,\dot{m})$ (arrived at by the authors of Ref. [7] through a combination of intuition and experience) and the diffusion map coordinates (Ψ_1,Ψ_2) (which were arrived at automatically based on the intrinsic geometry of the data). It is also interesting that the diffusion map representation appears to give well-filtered phase portraits, less sensitive to noise than the $(m,\dot{m})$ but also the principal component-based ones; this actually enhances the computational quantification of the underlying Hopf bifurcation through Poincaré map sections.

One might consider as a shortcoming of the diffusion map approach the fact that, as new particle simulation data become available, one does not have explicit formulas for

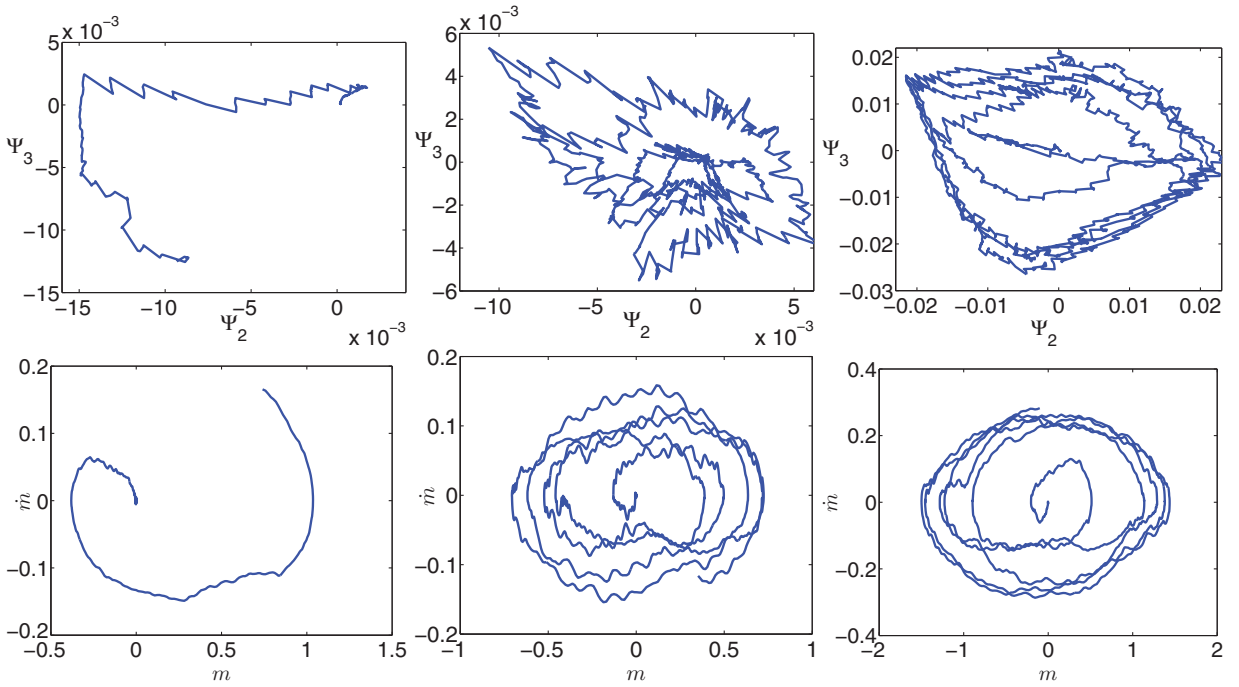

FIG. 8. (Color online) Phase space trajectories in diffusion map space (top), reconstructed through the Nyström extension [(6)] from new data points, compared to those in terms of the variables in Ref. [7] (bottom); transients are shown for the blocked regime $w = 0.53$ (left), just after the onset of oscillations for $w = 0.58$ (middle), and for $w = 0.65$ (right). For $w = 0.53$, the stationary state is stable and the trajectories approach this stable equilibrium; for $w = 0.58$ and $w = 0.65$ the trajectories spiral away from the unstable equilibrium towards the coarse limit cycle. Both embeddings exhibit the same qualitative behavior.

their embedding; in PCA, by contrast, new data on the same low-dimensional linear subspace can be expressed in PCA coordinates through a projection with a few simple inner products. This can be overcome through the use of the Nyström extension, as long as the new data points remain close to (have a sufficient number of close neighbors on) the regions of the low-dimensional manifold already sampled. Figure 8 illustrates this by presenting newly computed transients in Nyström-reconstructed diffusion map space (top) [and their counterparts in $(m,\dot{m})$ space, bottom] for $w = 0.53 < w^*$, $w = 0.58 > w^*$ and $w = 0.65 > w^*$. All trajectories were computed for 2000 time steps. The reference data set for the Nyström extensions through (6) consists of the data collected in the previous section over a range of door widths w. The plots in the variables m and $\dot{m}$ as well as the plots in Ψ_2 and Ψ_3 are clearly noisy; it appears that the Nyström extension, while allowing us to embed the data in diffusion map space, gives rise to trajectories comparably noisy to the m and $\dot{m}$ ones. Smoother trajectories could be obtained by using reference embeddings computed for a fixed door width. Nevertheless, both embeddings show the same qualitative behavior, namely a convergence to a fixed point for $w < w^*$ and oscillatory behavior for $w > w^*$.

E. Transition to oscillatory regime

In order to establish the Hopf nature of the bifurcation at $w^* = 0.55$ proposed in Ref. [7], the amplitude of Ψ_2 in Fig. 9 is plotted against the bifurcation parameter w. The figure clearly shows the gradual growth and subsequent saturation of the oscillation amplitude as a function of the door width; there is relatively little variation in the oscillation period (not shown). A comparison with the results from a direct simulation in the variables $(m,\dot{m})$ shows the same bifurcation point $w^* = 0.55$. A detailed bifurcation analysis is found in Ref. [7] and is not the focus of this paper.

Careful inspection of the snapshots of the oscillatory dynamics shows that the period of the oscillations is influenced by the return, through the periodic boundary conditions, of particles that passed through the door in the previous crossing surge. Small amplitude oscillations (close to the onset of the instability) correspond to fewer particles crossing in such a single surge event. If the particles did not have the opportunity to reenter, the problem would not be a stationary one—the density of particles congregating at each side of the door would gradually diminish after every surge, and this would

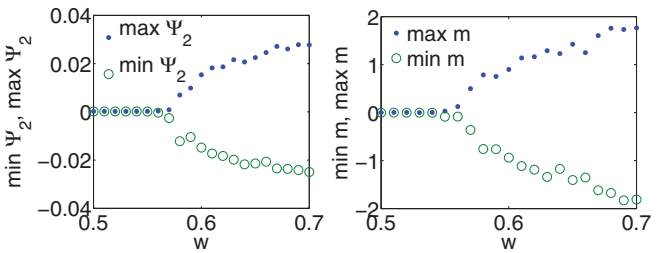

FIG. 9. (Color online) Bifurcation diagrams in diffusion map space ($\max \Psi_2, \min \Psi_2$) (left) and in the ($\max m, \min m$) variables of Ref. [7] (right). Both diagrams are consistent with an apparent Hopf bifurcation at a critical door width of $w^* = 0.55$. The panel in diffusion map space on the left appears slightly smoother.

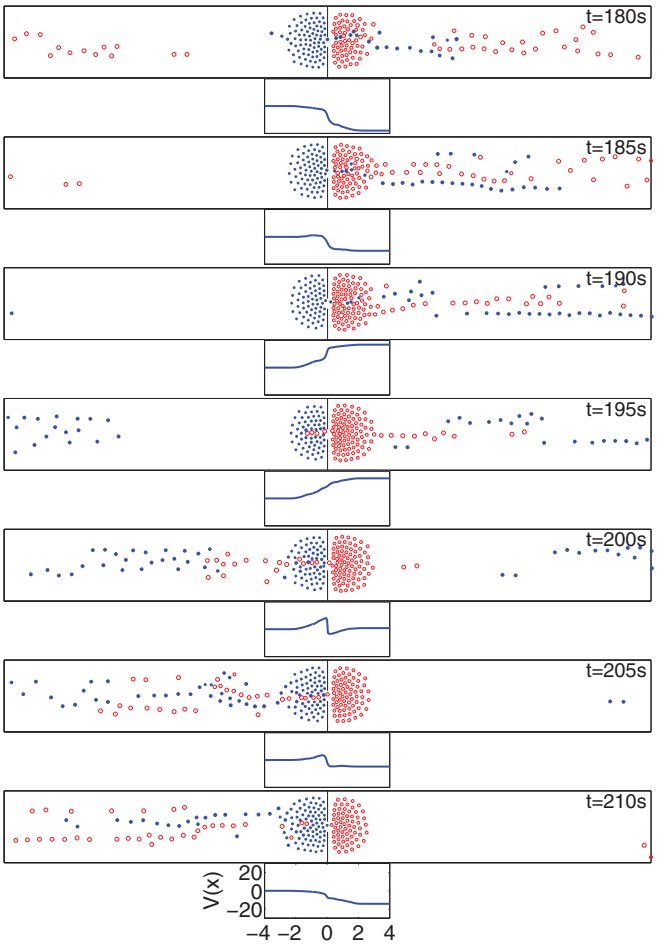

FIG. 10. (Color online) Potential $V(x)$ over x below the corresponding microscopic configuration for different times of the simulation with door width $w = 0.7$.

of course affect the switching times between crossings from alternate sides of the door. To provide an intuitive physical explanation of the mechanism underlying these switches, we show, in Fig. 10, a number of snapshots distributed along the oscillation; for each one of them we have calculated the force exerted on a colorless test pedestrian positioned on the corridor centerline by all the surrounding pedestrians. We then plot the integral of this force as a function of the position along the centerline, obtaining a sense of a pressure potential felt by the particle; the gradient of this computed quantity influences the particle motion. This potential is computed as follows: for any given fixed snapshot, the test pedestrian is located in a sequence of positions along the corridor axis, i.e., $y = 0$ and $x \in [-4,4]$ with $\Delta x = 0.05$; the force in corridor direction [(A2)] applied by the other pedestrians, $F_x(x)$, is computed, and the potential is then obtained as

$$V(x) = -\int_{-4}^{x} F_x(x')dx' \tag{13}$$

using the gauge $V(-4) = 0$. Figure 10 interleaves the evolution of the detailed state and the evolution of the potential over one oscillation period. We clearly see this pressure potential building up on the blue left-hand side of the door during the first couple of snapshots, and then reversing (as now the density of red circle particles on the red right-hand side of the door is larger). When, eventually, the blue dot particles that crossed return through the periodic boundary conditions, the blue dot particle density on the blue left-hand side of the door is roughly restored, and the pressure potential repeats.

The interactions of the pedestrians in front of the door clearly lead to surges through the door from alternate sides; the occurrence of these surges depends (as we tried to argue above) on the density of pedestrians in the neighborhood of the door. The periodic boundary conditions and the return of pedestrians through them back to the door is important in replenishing the particle density close to the door, making the behavior not just alternating, but regularly periodic and making the instability appear like a Hopf bifurcation. Other mechanisms of replenishing particle density, without periodic boundary conditions, such as the random injection of particles at some average rate, can also lead to regular periodic behavior by balancing overall particle inflow and outflow. On the other hand, keeping the densities close to the door effectively constant (through idealized particle reservoirs) would make the problem take a bistable switching aspect, with a noise-induced distribution of switching times.

In all computations in this paper, the number of red circle and blue dot particles, as well as their intrinsic properties, were taken to be identical. This makes the oscillation have a symmetric nature in time: evolving forward for half a period (coarsely) commutes with reflecting the corridor around the door ($x = 0$) and flipping particle colors (grayscales). Such a symmetry (picturesquely called ponies on a merry-go-round (POMs) [37,38]) has implications for the bifurcation scenarios possible [39]. Breaking any of these symmetries would destroy the POM nature of the limit cycles we observe.

V. CONCLUSION AND OUTLOOK

We have shown in this article that diffusion maps can be successfully applied to assist the study of pedestrian dynamics. The pedestrian model is used as a representative example of a particle system with time-scale separation. Not only does the use of diffusion maps avoid the need for user-specific selection of good coarse variables, it also appears here to lead to filtered, smoother coarse trajectories, which can be helpful in coarse-grained bifurcation analysis. Our study confirms that the use of diffusion maps is well suited for studying systems where intuition about good coarse variables is lacking.

On the other side of the coin, the lack of physical interpretation of the diffusion map embedding is a nontrivial drawback. There is clearly an upfront cost in the computation of diffusion maps, involving the computation of (many) pairwise distances and some large-scale linear algebra eigencomputations. As we already mentioned, careful data structuring, fast nearest-neighbor detection algorithms, and matrix sparsity can help make this task easier, so that the methods become applicable to very large data sets.

The automated discovery of good coarse variables is a crucial enabling technology for multiscale computations, and, in particular, for the application of equation-free techniques to new problems. In such computations, the ability to routinely transform from fine-scale (for us here, particle position and

velocity) space to coarse-grained (here diffusion map) space is an important component of the algorithms. And while, as we discussed, tools like the Nyström extension help in one direction of this transformation (the restriction to diffusion map space), the converse transformation (the lifting) is nontrivial and can involve sophisticated multiscale interpolation and extrapolation techniques (see for example the discussion in Ref. [40] and the techniques proposed there).

In our work we implicitly assumed that the appropriate macroscopic description would be in terms of deterministic, ordinary differential equations; indeed Figs. 4, 5, and 6 and our discussion of the Hopf nature of the underlying bifurcation support this assumption: our system is well modeled by a set of ODEs (modulo a little noise).

In other problems, however, stochastic effects may well be more pronounced, and the appropriate coarse description might be in terms of effective SDEs or the associated effective Fokker-Planck equations [41]; and while the estimation of coarse right-hand sides would now include both effective drifts and effective diffusivities, the role of diffusion maps remains the same: detect the coarse variables in terms of which the effective SDEs can be formulated.

It is worth noting that we have already mentioned a version of our pedestrian problem where such an effective stochastic reduction might be called for: the constant density case, where one might expect noise-induced bistable switching to lead to a distribution of surge times from alternate sides of the door.

While in this paper we reexamined, through a different approach, phenomena whose existence we already knew, our pedestrian model still possesses in its parameter space a wealth of possible behaviors that will pose their own challenges to data mining and coarse bifurcation computations. We know, for example, that for large enough door widths lane formation (the formation of striations in the pedestrian traffic pattern) will take place. The spatiotemporal nature of this instability will clearly need more and different coarse variables than the ones that were sufficient here, and the nature of the underlying coarse instability is yet to be explored.

Finally, as another interesting avenue of research, we mention the possibility of using coarse graining in the study of *heterogeneous* crowds. Here, we had blue dot and red circle pedestrians, that were otherwise identical. If the intrinsic properties of these particles (in the form of preferred target velocities, different interactions with the walls, different reaction times and/or perception distances, etc.) are not fixed, but sampled from a distribution, the problem acquires a new dimension. Recent developments involving mathematical tools from uncertainty quantification hold promise towards effective coarse graining in such heterogeneous problems [42,43].

ACKNOWLEDGMENTS

C.M. and J.S. would like to acknowledge the hospitality of Princeton University during research visits. C.M. would like to thank Otto Mønsted Fonden for financial support of the research stay in Princeton. J.S. would also like to thank the Danish Research Council FTP (09-065890/FTP) and the Villum Fonden (VKR-Centre of Excellence Ocean Life) for financial support. The work of I.G.K. and P.L. was partially supported by the US Department of Energy (DE-SC0005176 and DE-SC0002097).

APPENDIX A: SOCIAL FORCES AND MODEL PARAMETERS

The *inherent driving force* of a pedestrian i is given by its target direction $\mathbf{e}_i^0$ and velocity v^0. The force $\mathbf{F}_i^0$ tries to align the trajectory of the pedestrian to its target direction with

$$\mathbf{F}_i^0 = \tau^{-1}\left[v^0\mathbf{e}_i^0 - \dot{\mathbf{z}}_i(t)\right], \tag{A1}$$

where $\dot{\mathbf{z}}_i(t)$ is the velocity of pedestrian i at time t and τ is the reaction time.

The *pedestrian interaction* is modeled by the term

$$\begin{aligned}\mathbf{f}_{ij} &= \mathbf{f}(V,\sigma,\mathbf{r}_{ij}) \\ &= \begin{cases} -V\left[\tan\left(g(\|\mathbf{r}_{ij}\|)\right) - g(\|\mathbf{r}_{ij}\|)\right]\frac{\mathbf{r}_{ij}}{\|\mathbf{r}_{ij}\|}, & \|\mathbf{r}_{ij}\| \leqslant \sigma \\ 0, & \|\mathbf{r}_{ij}\| > \sigma. \end{cases}\end{aligned} \tag{A2}$$

Here, $\mathbf{r}_{ij}$ is the vector from pedestrian i to j, $\|\mathbf{r}_{ij}\|$ is the distance between pedestrians i and j, and V is the repulsion strength. The cutoff length σ reflects the fact that largely separated pedestrians do not influence each other. $g(\|\cdot\|) = \frac{\pi}{2}\left(\frac{\|\cdot\|}{\sigma} - 1\right)$ is introduced as a shorthand notation.

The third term describes the *interaction of pedestrians with walls*. Since pedestrians try to avoid collisions with walls, a repulsion force

$$\mathbf{f}_{iB} = \mathbf{f}(V_B,R,\mathbf{r}_{iB}) \tag{A3}$$

is introduced. It has the same functional dependence as (A2), but a different repulsion strength V_B and different range $R > \sigma$. The vector $\mathbf{r}_{iB}$ is the distance vector between pedestrian i and the closest point on boundary B.

The last term in (1) is the *additive noise*, which is introduced to avoid deadlock situations. It also reflects that people tend to avoid collisions with other persons by moving to one preferred side (assuming the pedestrian live in the same country). Here, we assume that the pedestrians tend to move to the right, leading to a noise term as

$$\mathbf{n}_i = n_i^{\parallel}\mathbf{e}_i + n_i^{\perp}\mathbf{e}_i^{\perp} = n_i^{\parallel}\mathbf{e}_i + n_i^{\perp}\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}\mathbf{e}_i, \tag{A4}$$

where $n_i^{\parallel}$ and $n_i^{\perp}$ are the noise components parallel and perpendicular to the target direction $\mathbf{e}_i$, respectively. Both components are normally distributed with

$$n_i^{\parallel} \sim \mathcal{N}(0,(s^{\parallel})^2) \quad \text{and} \quad n_i^{\perp} \sim \mathcal{N}(0.1s^{\perp},(s^{\perp})^2) \tag{A5}$$

where $s^{\parallel} = 0.00158$ and $s^{\perp} = 0.0632$. $\mathcal{N}(\mu,s^2)$ is a normal distribution with mean μ and variance s^2.

APPENDIX B: DIFFUSION MAP ALGORITHM

For convenience of the reader and definition of notation, the diffusion map algorithm is summarized in the following, see Refs. [21,22] for details. It has been successfully applied to physical problems, e.g., the description of a driven interface in an Ising variant model [26] or the study of the dynamics of animal groups [28].

Given a set of properly scaled observations

$$X = \{\mathbf{x}_i \in \mathbb{R}^n | i = 1, \ldots, N\}, \tag{B1}$$

where N is the number of observations and n is the dimension of the data, we compute all pairwise distances and arrange them in the distance matrix d:

$$d_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|, \quad i, j = 1, \ldots, N. \tag{B2}$$

Usually, $\|\cdot\|$ is the Euclidean norm in $\mathbb{R}^n$, but it can also be chosen to be problem specific [cf. (12), where we use a norm on moments]. The scaling parameter ϵ is used to compute the affinity matrix A as

$$A_{ij} = \exp\left(-\frac{d_{ij}^2}{\epsilon^2}\right). \tag{B3}$$

Note that large distances in d are being mapped to small pairwise affinities in A. The rows of A are normalized to obtain a Markov transition matrix

$$M_{ij} = \left(\sum_{j=1}^{N} A_{ij}\right)^{-1} A_{ij}, \tag{B4}$$

where M_{ij} is the one-step probability of transitioning from $\mathbf{x}_i$ to $\mathbf{x}_j$. M therefore defines a diffusion process, i.e., a Markov chain, on X.

Using M, the time-t diffusion distance D_t between two points $\mathbf{x}_i, \mathbf{x}_j \in X$ is defined as

$$D_t(\mathbf{x}_i, \mathbf{x}_j)^2 = \sum_k \frac{\left|M_{ik}^t - M_{jk}^t\right|^2}{\phi_0(k)}, \tag{B5}$$

where

$$\phi_0(k) = \frac{\sum_{j=1}^{N} A_{kj}}{\sum_{i=1}^{N}\sum_{j=1}^{N} A_{ij}} \tag{B6}$$

is the stationary distribution (cf. [44]). The diffusion distance D_t measures the difference in probability for transitioning from state i and j to state k, respectively. If the probabilities are the same, state k is equally well connected to states i and j by a diffusion process. (In a continuous analog, D_t measures the overlap of the two distributions after time t resulting from initializing with δ distributions centered at x_i and x_j, respectively.) The terms are normalized by the stationary distribution $\phi_0(k)$, which represents the probability to find the diffusion process in state k for $t \to \infty$. We then define a transformation from points $\mathbf{x}_i \in X$ to points $\mathbf{y}_i \in Y$ such that the Euclidean distance in Y equals the diffusion distance in X, i.e.,

$$\|\mathbf{y}_i - \mathbf{y}_j\|^2 = D_t(\mathbf{x}_i, \mathbf{x}_j)^2. \tag{B7}$$

A termwise comparison with (B5) using $p_t(\mathbf{x}_i, \mathbf{x}_k) = M_{ik}^t / \sqrt{\phi_0(k)}$ yields the transformed coordinates $\mathbf{y}$:

$$\mathbf{y}_i = [p_t(\mathbf{x}_i, \mathbf{x}_1), \ldots, p_t(\mathbf{x}_i, \mathbf{x}_N)]^T, \tag{B8}$$

where $p_t(\mathbf{x}_i, \mathbf{x}_j)$ is the probability of transitioning from point $\mathbf{x}_i$ to $\mathbf{x}_j$ in t steps. A useful approximation of this embedding can be formulated in terms of the eigenvalues and eigenvectors of M; if the data indeed lie on (or reasonably close to) a low-dimensional manifold, only a few (say, the first $k+1$) leading eigenvalues/eigenvectors need to be computed

$$M\boldsymbol{\Psi}_i = \lambda_i \boldsymbol{\Psi}_i, \quad i = 0, \ldots, k, \tag{B9}$$

where $\lambda_i > \lambda_j$ for $i < j$, assuming nondegeneracy of the eigenvalues. Since M is a Markov matrix, $\lambda_0 = 1$ and Ψ_0 is a vector containing only entries of 1. The transformation from data space to diffusion map space, with $\hat{\mathbf{y}}_i \in \mathbb{R}^k$, is given by

$$\mathbf{x}_i \mapsto \left[\lambda_1^t \boldsymbol{\Psi}_{1,i}, \ldots, \lambda_k^t \boldsymbol{\Psi}_{k,i}\right]^T = \hat{\mathbf{y}}_i, \tag{B10}$$

where $\boldsymbol{\Psi}_{u,v}$ is the vth component of eigenvector $\boldsymbol{\Psi}_u$. For $k \ll N$ this transformation constitutes a dimension reduction scheme.

Note that for larger t (i.e., when using longer times in the diffusion process) higher eigenvalues become increasingly damped as we take their powers; fewer leading (slower) components of $\hat{\mathbf{y}}_i$ will then suffice to approximate the data to a given accuracy. This allows the structure of the data manifold to be investigated at different scales [45]. This approach, through its ability to parametrize curved, nonlinear, manifolds can be more economical in the representation of data possessing such structure than linear methods (such as PCA or MDS).

[1] D. Helbing and P. Molnár, Phys. Rev. E **51**, 4282 (1995).
[2] D. Helbing, Rev. Mod. Phys. **73**, 1067 (2001).
[3] D. Helbing, I. Farkas, and T. Vicsek, Nature (London) **407**, 487 (2000).
[4] D. Helbing, R. Jiang, and M. Treiber, Phys. Rev. E **72**, 046130 (2005).
[5] N. Bellomo and C. Dogbe, SIAM Rev. **53**, 409 (2011).
[6] A. J. Bernoff and C. M. Topaz, SIAM J. Appl. Dyn. Syst. **10**, 212 (2011).
[7] O. Corradi, P. Hjorth, and J. Starke, SIAM J. Appl. Dyn. Syst. **11**, 1007 (2012).
[8] N. Fenichel, J. Differ. Eqn. **31**, 53 (1979).
[9] H. Haken, in *Advanced Synergetics: Instability Hierarchies of Self-Organizing Systems and Devices*, Springer Series in Synergetics 20 (Springer, Berlin, 1983), pp. 1–356.
[10] H. Haken, in *Synergetics: An Introduction. Nonequilibrium Phase Transitions and Self Organization in Physics, Chemistry and Biology*, 3rd ed., Springer Series in Synergetics 1 (Springer, Berlin, 1983), pp. 1–371.
[11] I. G. Kevrekidis, C. W. Gear, J. M. Hyman, P. G. Kevrekidis, O. Runborg, and C. Theodoropoulos, Comm. Math. Sci. **1**, 715 (2003).
[12] I. G. Kevrekidis and G. Samaey, Annu. Rev. Phys. Chem. **60**, 321 (2009).
[13] K. Pearson, Philosophical Magazine Series **6**, 559 (1901).
[14] W. S. Torgerson, Psychometrika **17**, 401 (1952).
[15] T. Cox and M. Cox, *Multidimensional Scaling* (Chapman & Hall, London, 1994).
[16] P. Holmes, J. Lumley, and G. Berkooz, in *Turbulence, Coherent Structures, Dynamical Systems and Symmetry* (Cambridge University Press, Cambridge, 1998), pp. 1–420.

[17] A. Talwalkar, S. Kumar, and H. Rowley, in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Anchorage, AK, 2008* (IEEE, Piscataway, NJ, 2008), pp. 1–8.
[18] S. T. Roweis and L. K. Saul, Science **290**, 2323 (2000).
[19] D. Yan, L. Huang, and M. I. Jordan, in *KDD '09: Proceedings of the 15th ACM SIGKDD International Conference on Knowledge, Discovery, and Data Mining* (ACM, New York, 2009), pp. 907–916.
[20] M. Belkin and P. Niyogi, in *Advances in Neural Information Processing Systems 14* (MIT Press, Cambridge, 2002), pp. 585–591.
[21] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker, Proc. Natl. Acad. Sci. USA **102**, 7426 (2005).
[22] R. R. Coifman and S. Lafon, Applied and Computational Harmonic Analysis **21**, 5 (2006).
[23] M. Belkin and P. Niyogi, Neural Comput. **15**, 1373 (2003).
[24] A. Kolpas, J. Moehlis, T. A. Frewen, and I. G. Kevrekidis, Math. Biosci. **214**, 49 (2008).
[25] R. R. Coifman, I. G. Kevrekidis, S. Lafon, M. Maggioni, and B. Nadler, Multiscale Modeling and Simulation **7**, 842 (2008).
[26] B. E. Sonday, M. Haataja, and I. G. Kevrekidis, Phys. Rev. E **80**, 031102 (2009).
[27] C. Laing, T. Frewen, and I. Kevrekidis, J. Comput. Neurosci. **28**, 459 (2010).
[28] T. A. Frewen, I. D. Couzin, A. Kolpas, J. Moehlis, R. Coifman, and I. G. Kevrekidis, in *Coping with Complexity: Model Reduction and Data Analysis*, Lecture Notes in Computational Science and Engineering, edited by A. N. Gorban, D. Roose, T. J. Barth, M. Griebel, D. E. Keyes, R. M. Nieminen, D. Roose, and T. Schlick, Vol. 75 (Springer, Berlin, 2011), pp. 299–309.
[29] M. A. Rohrdanz, W. Zheng, M. Maggioni, and C. Clementi, J. Chem. Phys. **134**, 124116 (2011).
[30] A. L. Ferguson, A. Z. Panagiotopoulos, I. G. Kevrekidis, and P. G. Debenedetti, Chem. Phys. Lett. **509**, 1 (2011).
[31] T. Berry, J. Cressman, Z. Gregurić-Ferenček, and T. Sauer, SIAM J. Appl. Dyn. Syst. **12**, 618 (2013).
[32] Isomap, http://isomap.stanford.edu/
[33] P. Grassberger and I. Procaccia, Phys. Nonlinear Phenom. **9**, 189 (1983).
[34] S. Kumar, M. Mohri, and A. Talwalkar, J. Mach. Learn. Res. **13**, 981 (2012).
[35] Y. Rubner, C. Tomasi, and L. Guibas, Int. J. Comput. Vis. **40**, 99 (2000).
[36] CPU: Intel Xeon W3690 3.47GHz ×6, Memory: 6GB.
[37] D. G. Aronson, M. Golubitsky, and J. Mallet-Paret, Nonlinearity **4**, 903 (1991).
[38] D. G. Aronson, M. Golubitsky, and M. Krupa, Nonlinearity **4**, 861 (1991).
[39] J. W. Swift and K. Wiesenfeld, Phys. Rev. Lett. **52**, 705 (1984).
[40] E. Chiavazzo, C. W. Gear, C. J. Dsilva, N. Rabin, and I. G. Kevrekidis, arXiv:1307.6849 [math.DS] (2013).
[41] B. Nadler, S. Lafon, R. R. Coifman, and I. G. Kevrekidis, Applied and Computational Harmonic Analysis **21**, 113 (2006).
[42] S. J. Moon, R. Ghanem, and I. G. Kevrekidis, Phys. Rev. Lett. **96**, 144101 (2006).
[43] C. Laing, Y. Zou, B. Smith, and I. Kevrekidis, J. Math. Neurosci. **2**, 1 (2012).
[44] S. Lafon and A. Lee, IEEE Trans. Pattern Anal. Mach. Intell. **28**, 1393 (2006).
[45] J. de la Porte, B. Herbst, W. Hereman, and S. van der Walt, *Nineteenth Annual Symposium of the Pattern Recognition Association of South Africa (PRASA)* (Pattern Recognition Association of South Africa, Johannesburg, 2008).