

Parameter estimation from aggregate observations: a Wasserstein distance-based sequential Monte Carlo sampler

Cheng, Chen; Wen, Linjie; Li, Jinglai

DOI:
[10.1098/rsos.230275](https://doi.org/10.1098/rsos.230275)

License:
Creative Commons: Attribution (CC BY)

Document Version
Publisher's PDF, also known as Version of record

Citation for published version (Harvard):
Cheng, C, Wen, L & Li, J 2023, 'Parameter estimation from aggregate observations: a Wasserstein distance-based sequential Monte Carlo sampler', *Royal Society Open Science*, vol. 10, no. 8.
<https://doi.org/10.1098/rsos.230275>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Research



Cite this article: Cheng C, Wen L, Li J. 2023
Parameter estimation from aggregate
observations: a Wasserstein distance-based
sequential Monte Carlo sampler. *R. Soc. Open Sci.*
10: 230275.
<https://doi.org/10.1098/rsos.230275>

Received: 9 March 2023

Accepted: 21 June 2023

Subject Category:

Mathematics

Subject Areas:

statistics/mathematical modelling/applied
mathematics

Keywords:

parameter estimation, sequential Monte Carlo
sampler, likelihood-free inference, Wasserstein
distance

Author for correspondence:

Jinglai Li

e-mail: j.li.10@bham.ac.uk

Parameter estimation from aggregate observations: a Wasserstein distance-based sequential Monte Carlo sampler

Chen Cheng¹, Linjie Wen² and Jinglai Li³

¹School of Mathematical Sciences, Shanghai Jiao Tong University, Shanghai 200240, People's Republic of China

²School of Earth and Space Sciences, Peking University, 5 Yiheyuan Rd, Beijing 100871, People's Republic of China

³School of Mathematics, University of Birmingham, Birmingham B15 2TT, UK

CC, 0000-0001-9773-2993

In this work, we study systems consisting of a group of moving particles. In such systems, often some important parameters are unknown and have to be estimated from observed data. Such parameter estimation problems can often be solved via a Bayesian inference framework. However, in many practical problems, only data at the aggregate level is available and as a result the likelihood function is not available, which poses a challenge for Bayesian methods. In particular, we consider the situation where the distributions of the particles are observed. We propose a Wasserstein distance (WD)-based sequential Monte Carlo sampler to solve the problem: the WD is used to measure the similarity between the observed and the simulated particle distributions and the sequential Monte Carlo samplers is used to deal with the sequentially available observations. Two real-world examples are provided to demonstrate the performance of the proposed method.

1. Introduction

In transportation science, it is often of critical importance to study the collective behaviour of a group of moving objectives (referred to as particles hereafter). A good example of such problems is the dynamics of a crowd of human pedestrians, which has important applications in urban planning and safety management [1,2]. Other examples of such systems include traffic flows [3], swarm robots [4] and so on. Modelling these collective behaviours has attracted considerable attention in multiple disciplines, and




Figure 1. An illustrative example of the TV distance and the Wasserstein distance where the histograms of three probability measures, U , V and W , are shown. The TV distance between U and V and that between U and W are the same. The WD between U and V is 0.03, while the WD between U and W is 0.3. See the main text for more information.

various models have been proposed in the past decades. Despite the modelling advances, some issues are yet to be solved. In particular, certain key information in the system may be missing, and as a result some important model parameters are not known to the practitioners. Examples of such a situation may appear in pedestrian crowds [5] and traffic flows [6,7]. In practice, these model parameters are often estimated by fitting the real-world observation data into the mathematical models, often via a Bayesian framework.

The Bayesian parameter estimation is conceptually straightforward and has been used to estimate model parameters in many related problems [8–11]. However, it can be highly challenging to apply the Bayesian method to the microscopic models of such systems—i.e. models that directly describe the dynamics of the individual particles in the system. A main difficulty is that, in order to conduct the standard Bayesian inference for a microscopic model, one needs to track each particle, which is extremely difficult or even impossible when the ensemble size is large. Instead, in reality, one often observes the aggregate data that are collected at the ensemble level and do not characterize the state of each individual particle. As will be explained later, determining a suitable distribution for aggregate measures in this case can be challenging, potentially resulting in the unavailability of an analytically derived likelihood function. Additionally, the presence of unknown measurement noise further complicates the tractability of the likelihood. As a consequence, standard posterior computation methods such as Markov Chain Monte Carlo (MCMC) cannot be used. A large class of approximate Bayesian inference techniques have been developed to address such likelihood-free problems, such as the approximate Bayesian computation (ABC) [12] and other methods, e.g. [13]. These techniques have been used to conduct likelihood-free inference from aggregate observations. For instance, macroscopic observation data measuring the flow is used to calibrate a microscopic model for crowd dynamics using ABC methods [10].

In the present work, we consider a special case of aggregate observations—namely, it is the distribution of the particles that is observed. Our main purpose is to develop a generic method that can compute the posterior in such problems, by taking advantage of the rich literature of the likelihood-free inference. As in our problem of interest, usually the data are observed in a sequential manner. We choose to base our method on the sequential Monte Carlo sampler (SMCS) [14], an extension of the particle filter, for its ability to deal with sequentially available data. Some variants of SMCS can deal with likelihood-free inference [15] and these methods typically rely on the distance (e.g. the Euclidean distance) between the simulated data and the observed data. To this end, a similar problem has been considered in the crowd dynamics context in [11], where a crowd flow forecasting method based on particle filter is proposed to estimate both the crowd state and latent parameters from the aggregate density observation data. In [11], the total-variation (TV) distance is used to measure the similarity between the observed and the simulated distributions. However, such distance measures may become ineffective when observation noise is present: in figure 1, we provide an example in which the TV distance between U and V and that between U and W are the same, which makes it unsuited for measuring the distance between the observed and the simulated distributions. To address the issue, we propose to use the Wasserstein distance (WD) [16] as a distance measure between the simulated and the observed distributions. The WD is a commonly used measure for similarity between two distributions, which loosely speaking, is the minimal cost for transforming one distribution into the other. In figure 1, the WD between U and V is computed as 0.03 while the WD between U and W is determined as 0.3. This discrepancy more reasonably characterizes the differences

between these distributions and addresses the aforementioned issue with the TV distance. WD has been widely used in many applications in statistics and machine learning, such as [17,18], as it has many desired properties. Additionally, it has been shown to be effective in mitigating information loss caused by the utilization of summary statistics in basic ABC techniques [19]. In this work we develop a WD-based SMCS method which can deal with the aforementioned problems where the observations are particle distributions (referred to as the aggregate data hereafter).

The rest of the work is organized as follows. We first discuss the likelihood-free sequential inference problem in §2, and present the proposed WD-based SMCS in §3. Numerical examples are provided in §4 to demonstrate the performance of the proposed method. Finally, §5 offers some conclusions.

2. Likelihood-free sequential Bayesian inference

2.1. Problem set-up

We start with a generic set-up of the sequential Bayesian inference problem considered here. Suppose that we have measurements y_1, \dots, y_T (with $y_t \in \mathbb{R}^{n_y}$ for $t = 1, \dots, T$) that are collected sequentially in time from a dynamical system model, and from the data we want to estimate the unknown parameter $\theta \in \mathbb{R}^{n_\theta}$. In such problems, after collecting t data points y_1, \dots, y_t , we can compute the posterior distribution of θ conditional on measurements y_1, \dots, y_t in the following form (assuming all the measurements are independent conditional on θ):

$$\pi_t(\theta) = \pi(\theta|y_1, \dots, y_t) \propto \pi_0(\theta) \prod_{i=1}^t \pi(y_i|\theta), \quad (2.1)$$

where $\pi_0(\theta)$ is the prior distribution of θ that is usually specified by the users, or equivalently,

$$\pi_t(\theta) \propto \pi_{t-1}(\theta) \pi(y_t|\theta). \quad (2.2)$$

Our goal is to compute $\pi_T(\theta)$, i.e. the final posterior distribution after collecting all data points y_1, \dots, y_T .

2.2. Sequential Monte Carlo sampler

Since the posterior distribution is typically not analytically available, a common practice is to characterize the distribution by drawing samples or realizations from it. In principle, the posterior $\pi_T(\theta)$ can be sampled with the standard methods such as MCMC. However, as is pointed out in [20,21], the posterior in equation (2.1) poses challenges for usual MCMC methods especially when T is large, as they need to query the likelihood function $\pi(\cdot|\theta)$ a large number of times. To this end, the SMCS method is particularly designed to exploit the sequential structure of such problems.

SMCS is a sequential importance sampling method, and as such it generates weighted samples from the posterior distribution, where the samples represent possible values from the posterior, while the weights assign importance to each sample based on their contribution to capturing characteristics of the distribution. We here describe SMCS in a recursive formulation. Given an arbitrary conditional distribution $L_{t-1}(\theta_{t-1}|\theta_t)$, one can construct the following joint distribution of θ_{t-1} and θ_t :

$$p_t(\theta_{t-1}, \theta_t) = \pi_t(\theta_t) L_{t-1}(\theta_{t-1}|\theta_t), \quad (2.3)$$

whose marginal distribution over θ_{t-1} is $\pi_t(\theta_t)$. Importance sampling [22] is applied to draw samples for $p_t(\theta_{t-1}, \theta_t)$, where the proposal distribution $q_t(\theta_{t-1}, \theta_t)$ is constructed in the form of

$$q_t(\theta_{t-1}, \theta_t) = q_{t-1}(\theta_{t-1}) K_t(\theta_t|\theta_{t-1}), \quad (2.4)$$

from which we can directly draw samples (θ_{t-1}, θ_t) . Here $q_{t-1}(\theta_{t-1})$ is a marginal distribution and $K_t(\theta_t|\theta_{t-1})$ a conditional distribution. The weights of samples are computed according to

$$\begin{aligned} \omega_t(\theta_{t-1:t}) &= \frac{p_t(\theta_{t-1}, \theta_t)}{q_t(\theta_{t-1}, \theta_t)} = \frac{\pi_t(\theta_t) L_{t-1}(\theta_{t-1}|\theta_t)}{q_{t-1}(\theta_{t-1}) K_t(\theta_t|\theta_{t-1})} \\ &= \omega_{t-1}(\theta_{t-1}) \alpha_t(\theta_{t-1}, \theta_t), \end{aligned} \quad (2.5a)$$

where

$$\omega_{t-1}(\theta_{t-1}) = \frac{\pi_{t-1}(\theta_{t-1})}{q_{t-1}(\theta_{t-1})} \quad (2.5b)$$

is the weight at $t - 1$ step, and

$$\alpha_t(\theta_{t-1}, \theta_t) = \frac{\pi_t(\theta_t)L_{t-1}(\theta_{t-1}|\theta_t)}{\pi_{t-1}(\theta_{t-1})K_t(\theta_t|\theta_{t-1})}, \quad (2.5c)$$

is the incremental weight at the t -th step. Now that we have joint samples $\{(\theta_{t-1}^j, \theta_t^j), \omega_t^j\}_{j=1}^N$ from the joint distribution $p_t(\theta_{t-1}, \theta_t)$, it follows that samples $\{\theta_t^j, \omega_t^j\}_{j=1}^N$ correspond to the marginal distribution $\pi_t(\theta_t)$. The complete SMCS algorithm proceeds as follows:

- (i) let $t = 0$, draw samples $\{\theta_0^j\}_{j=1}^N$ from $q_0(\cdot)$, and compute $\omega_0^j = \pi_0(\theta_0^j)/q_0(\theta_0^j)$ for $j = 1, \dots, N$;
- (ii) let $t = t + 1$;
- (iii) draw θ_t^j from $K_t(\cdot|\theta_{t-1}^j)$ for each $j = 1, \dots, N$;
- (iv) compute ω_t^j using equation (2.5) for each $j = 1, \dots, N$;
- (v) go to step (ii) if $t < T$.

Finally the procedure returns a set of weighted samples $\{\theta_T^j, \omega_T^j\}_{j=1}^N$ from the posterior $\pi_T(\theta)$. Note that in SMCS algorithms, a resampling step is commonly used to alleviate the ‘sample degeneracy’ issue [14]. When the effective sample size (ESS) falls below a specific threshold (usually less than half the total number of samples), resampling is performed on the proposed samples to mitigate this issue.

From the discussions above, we can see that in SMCS the forward kernel $K_t(\theta_t|\theta_{t-1})$ and the backward kernel $L_{t-1}(\theta_{t-1}|\theta_t)$ are critical for the performance of the method and therefore have to be chosen carefully. Here, we adopt the MCMC moves developed in [14]. The sampling step (iii) in the SMCS algorithm with this forward kernel is constructed as follows. First, we choose a proposal distribution $k(\theta_t|\theta_{t-1})$ and draw a sample θ_t^* from it with a sample θ_{t-1}^j from the previous iteration. Next the sample θ_t^* is accepted or rejected according to the following acceptance probability:

$$a_t(\theta_t^*|\theta_{t-1}^j) = \min\left\{\frac{\pi_t(\theta_t^*)k(\theta_{t-1}^j|\theta_t^*)}{\pi_t(\theta_{t-1}^j)k(\theta_t^*|\theta_{t-1}^j)}, 1\right\}. \quad (2.6)$$

With this forward kernel, an approximate optimal backward kernel can be derived as

$$L_{t-1}(\theta_{t-1}|\theta_t) = \frac{\pi_t(\theta_{t-1})K_t(\theta_t|\theta_{t-1})}{\pi_t(\theta_t)}, \quad (2.7)$$

where the detailed derivation can be found in [14] and is not repeated here. With the choice of this approximate backward kernel, the incremental weight function $\alpha_t(\theta_{t-1}, \theta_t)$ in equation (2.5c) reduces to

$$\alpha_t(\theta_{t-1}, \theta_t) = \frac{\pi_t(\theta_{t-1})}{\pi_{t-1}(\theta_{t-1})}, \quad (2.8)$$

and when applied to the posterior distribution in equation (2.2), it becomes,

$$\alpha_t(\theta_{t-1}, \theta_t) = \pi(y_t|\theta_{t-1}). \quad (2.9)$$

Finally, we provide some remarks on the SMCS method:

- First we reinstate that SMCS is a special type of importance sampling (IS) approach, as it draws samples from a proposal distribution instead of the actual posterior, while correcting for the bias by assigning proper weights to each samples. Therefore, just like the standard IS, the method produces a set of sample-weight pairs which follow the posterior distribution. We refer to [22] for more details of IS. The main reason to use such a method is that it is challenging to draw samples directly from the posterior. SMCS is special as it constructs the joint target distribution as well as the joint proposal distribution in a special way, using the forward and backward kernel functions.
- The functions $K_t(\theta_t|\theta_{t-1})$ and $L_{t-1}(\theta_{t-1}|\theta_t)$ are both conditional distributions used within the SMCS method. They are not part of the inference problem and should not be regarded as likelihood functions. Instead, they can be understood as some auxiliary functions required by the SMCS method. More precisely they are used in constructing joint distributions along with a marginal distribution, and the terms ‘forward’ and ‘backward’ are associated with the dependence of their arguments.

2.3. Likelihood-free SMCS

It is easy to see that, to apply SMCS to the sequential inference problem described in §2.1, one needs the knowledge of the likelihood function $\pi(\cdot | \theta)$. In many real-world applications, including the problems of our interest in this paper, often the likelihood function is not explicitly available, and instead there exists a simulation model that can generate synthetic observation data that distribute according to the likelihood function $\pi(\cdot | \theta)$, given a parameter value θ . In this case the approximate (or likelihood-free) inference methods can be used. While noting that many such approximate inference methods have been proposed, we focus on the one based on SMCS, which was modified from [23].

Recall that, in the SMCS algorithm, the likelihood function is evaluated on two occasions: the first is to calculate the acceptance probability in equation (2.6) and the second is to update the sample weights in equation (2.9). The main idea of the approximate method is rather simple: one first generates a synthetic observation data \hat{y} from the simulation model and approximates the value of the likelihood function with a kernel function $H(\cdot, \cdot)$ that characterizes the similarity between \hat{y} and the actual data y :

$$\pi(y|\theta) \approx H(\hat{y}, y), \quad (2.10)$$

where \hat{y} is the synthetic observation data generated from $\pi(\hat{y}|\theta)$. We restate here that the kernel function H is chosen to characterize the similarity between \hat{y} and y , and as such, the more similar \hat{y} and y are, the higher the likelihood function value is.

3. Aggregate observations and the Wasserstein distance

The application problems of our interest are a special case of the sequential inference problems described in §2. Specifically, we consider the dynamics of an ensemble of n particles (such as pedestrians or vehicles), which is described by a discrete-time dynamical system

$$x_{t+1} = F_t(x_t; \theta), \quad (3.1)$$

where $x_t = (x_{1t}, \dots, x_{nt})$ with x_{it} (for $i = 1, \dots, n$) being the position of the i -th particle at time t , and θ is the unknown model parameter that we want to estimate.

In the standard set-up of such problems, one is able to observe the particle position x_t at each time t , from which the parameter θ is inferred, and in this case, the likelihood function is usually available. However, in many real-world problems, especially when the number of particles is large, it is not possible to track and locate each particle, and instead it is often much easier to observe how the particles are distributed at a given time t , often referred to as *aggregate observations*. Namely, at each $t = 1, \dots, T$, we observe the distribution of x_t , denoted as $\rho_t(x)$, and we aim to infer parameter θ from $\rho_t(x)$. It should be clear that, since the relation between the parameter and the aggregate observations are complex, it is usually challenging to derive the likelihood function analytically. As such likelihood-free methods need to be used for such problems. We adopt the likelihood-free SMCS method described in §2.3 to solve the problem, and recall that, in this framework, the key issue is to choose the kernel function $H(\cdot, \cdot)$ that measures the similarity between two data points that are essentially two distributions in the present problem. To this end, we choose to define the kernel function H via the WD of two distributions. The WD, also known as the earth mover's distance (EMD), is a popular choice for measuring dissimilarity or distance between two distributions, and widely used in real-world applications. Intuitively speaking the WD is calculated as the minimum 'cost' of transforming one distribution into the other. In what follows we discuss how to calculate the WD between two discrete-valued distributions. Let $\rho_U = \{(u_1, p_1^u), \dots, (u_m, p_m^u)\}$, $\rho_V = \{(v_1, p_1^v), \dots, (v_n, p_n^v)\}$ be two discrete-valued distributions, where p_i^u represents the density at the data point u_i . A function $c(u_i, v_j)$ is further defined as the cost for transporting a unit from $u_i \in U$ to $v_j \in V$, which is taken to be the Euclidean distance between u_i and v_j in this work. To calculate the WD, we first need to compute a 'transport map' $\gamma(u_i, v_j)$ that minimize the overall cost

$$\sum_{i=1}^m \sum_{j=1}^n c(u_i, v_j) \gamma(u_i, v_j), \quad (3.2)$$

subject to the following constraints:

$$\left. \begin{aligned} \sum_{j=1}^n \gamma(u_i, v_j) &= p_i^u, \\ \sum_{i=1}^m \gamma(u_i, v_j) &= p_j^v, \\ \gamma &\geq 0. \end{aligned} \right\} \quad (3.3)$$

and

Once the optimal transport map γ is found, the WD is defined as

$$D_W(\rho_U, \rho_V) = \frac{\sum_{i=1}^m \sum_{j=1}^n c(u_i, v_j) \gamma(u_i, v_j)}{\sum_{i=1}^m \sum_{j=1}^n \gamma(u_i, v_j)}. \quad (3.4)$$

For more details of the WD, we refer to [16] and the references therein. It follows that the kernel function can be defined as

$$H(\rho_U, \rho_V) = \frac{1}{\sqrt{2\pi}h} \exp\left(-\frac{D_W(\rho_U, \rho_V)^2}{2h^2}\right), \quad (3.5)$$

where h is the bandwidth parameter of the kernel. It is of essential importance to choose an appropriate value for h as it directly affects the sensitivity of the kernel to the WD value. The proper choice of h is usually problem dependent, and more importantly, for a particular problem, as the WDs of samples can vary across iterations, a fixed value of h may not be suitable for all the iterations. As a result, we adopt an adaptive approach where h is dynamically determined for each iteration. Namely, we set h to be the median WD among the samples. This adaptive strategy allows us to adjust the value of h according to varying WDs encountered during the estimation process, ensuring an appropriate level of sensitivity for the kernel. Finally we present the complete algorithm of the WD-based SMCS in algorithm 1.

Algorithm 1. Wasserstein distance-based sequential Monte Carlo sampler (WD-SMCS).

Input: Number of samples N , prior of parameters $\pi_0(\theta)$, observed distributions $y_{1:T}$, resampling threshold ESS_{th}

- 1: draw N samples $\{\theta_0^j\}_{j=1}^N$ from prior $\pi_0(\theta)$ and initialize weights $\omega_0^j = 1/N$;
- 2: **for** $t = 1, \dots, T$ **do**
- 3: draw candidate samples $\theta_t^{j*} \sim k(\cdot | \theta_{t-1}^j)$;
- 4: generate synthetic distributions $\hat{y}_{1:t}^j$ and $\hat{y}_{1:t}^{j*}$ with model parameter θ_{t-1}^j and θ_t^{j*} ;
- 5: compute pair-wise Wasserstein distance $D_W(\hat{y}_{1:t}^j, y_{1:t})$ and $D_W(\hat{y}_{1:t}^{j*}, y_{1:t})$;
- 6: compute approximated likelihood $\pi(y_{1:t} | \theta_t^{j*})$ and $\pi(y_{1:t} | \theta_{t-1}^j)$ via equations (2.10) and (3.5);
- 7: compute acceptance

$$a(\theta_t^{j*}, \theta_{t-1}^j) = \min\left\{\frac{\pi_t(\theta_t^{j*})k(\theta_{t-1}^j | \theta_t^{j*})}{\pi_t(\theta_{t-1}^j)k(\theta_t^{j*} | \theta_{t-1}^j)}, 1\right\};$$

- 8: generate random $u \sim U(0,1)$, if $u < a(\theta_t^{j*}, \theta_{t-1}^j)$, $\theta_t^j = \theta_t^{j*}$, else $\theta_t^j = \theta_{t-1}^j$;
- 9: update the weights with normalization

$$\omega_t^j = \omega_{t-1}^j \pi(y_t | \theta_t^j);$$

- 10: calculate $\text{ESS} = 1/\sum_{j=1}^N (\omega_t^j)^2$ and resample if $\text{ESS} < \text{ESS}_{th}$.
 - 11: **end for**
-

4. Numerical experiments

In this section, we consider two application problems, with which we conduct numerical experiments to examine the performance of the proposed method. The first example is the classic social force model (SFM) for the crowd dynamics of pedestrians. The second is the intelligent driver model (IDM) for the collective behaviour of traffic flow. We discuss the model details and the setup of the simulation scenarios in the following sections.




Figure 2. Schematic plots of the simulation scenarios. (Left) Collective escape towards a single exit. Pedestrians are depicted by dots and walls (i.e. obstacles) are depicted by lines. All pedestrians are moving towards the centre of the exit. (Right) Traffic flow on a four-lane highway. Vehicles are depicted by blue rectangles and white arrows on the roads indicate the driving direction of the vehicles.

4.1. The social force model

Our first example is the dynamics of pedestrian crowds, and in particular we consider a typical scenario of collective escape towards a single exit (see figure 2, left). We adopt the SFM [24] to describe the crowd behaviour. Simply speaking SFM assumes a mixture of physical and socio-psychological forces influencing the crowd behaviour by considering personal motivations and environmental constraints. In this model, each pedestrian i of mass m_i and velocity v_i tends to move by a desired speed v_i^p along a certain direction e_i^p during the acceleration time τ_i . The resulting personal desire force F_p is

$$F_p = m_i \frac{v_i^p e_i^p - v_i}{\tau_i}. \quad (4.1)$$

Furthermore, pedestrians psychologically tend to keep a social distance between each other and avoid hitting walls. This is modelled by introducing ‘interaction force’ f_{ij} between pedestrians i and j and f_{iW} between pedestrian i and the wall, respectively. The total interaction force is

$$F_{\text{int}} = \sum_{j(\neq i)} f_{ij} + \sum_W f_{iW}. \quad (4.2)$$

Combining equations (4.1) and (4.2), we obtain the acceleration equation

$$m_i \frac{dv_i}{dt} = m_i \frac{v_i^p(t) e_i^p(t) - v_i(t)}{\tau_i} + \sum_{j(\neq i)} f_{ij} + \sum_W f_{iW}. \quad (4.3)$$

The position vector $r_i(t)$ is updated by the velocity $v_i(t) = dr_i/dt$.

The interaction force f_{ij} between pedestrian i and j is specified as follows. With the distance $d_{ij} = \|r_i - r_j\|$ between the two pedestrians’ centres of mass, the psychological tendency of pedestrian i to stay away from pedestrian j is described by a repulsive interaction force $A_i \exp[(r_{ij} - d_{ij})/B_i] n_{ij}$, where A_i and B_i are constants, indicating the strength and the range of the interaction, and $n_{ij} = (n_{ij}^1, n_{ij}^2) = (r_i - r_j)/d_{ij}$ is the normalized directional vector pointing from pedestrian j to i . The pedestrians touch each other if their distance d_{ij} is smaller than the sum $r_{ij} = r_i + r_j$ of their radius r_i and r_j . In our model, we specify a uniform value for the size of each pedestrian (table 1) [24]. Inspired by granular interactions, two additional forces are included in the model, which are essential for understanding the particular effects in panicking crowds: a ‘body force’ $k(r_{ij} - d_{ij}) n_{ij}$ counteracting body compression and a ‘sliding friction force’ $\kappa(r_{ij} - d_{ij}) \Delta v_{ij}^t t_{ij}$ impeding relative tangential motion, if pedestrians i and j are close enough. Here, $t_{ij} = (-n_{ij}^2, n_{ij}^1)$ means the tangential direction and $\Delta v_{ij}^t = (v_j - v_i) \cdot t_{ij}$ the tangential velocity difference, while k and κ are large constants, representing the bump and the friction effect. In summary, the interaction force f_{ij} between pedestrians i and j is given by

$$f_{ij} = \left\{ A_i \exp \left[\frac{(r_{ij} - d_{ij})}{B_i} \right] + k \mathbf{I}(r_{ij} - d_{ij}) \right\} n_{ij} + \kappa \mathbf{I}(r_{ij} - d_{ij}) \Delta v_{ij}^t t_{ij}, \quad (4.4)$$

where the indicator function $\mathbf{I}(r_{ij} - d_{ij})$ is zero for $r_{ij} - d_{ij} < 0$ and it is equal to $r_{ij} - d_{ij}$ otherwise.

Table 1. List of parameter values in SFM and the three parameters to infer are in bold.

variable	value	description
m	80 kg	mass of pedestrians
v^p	1.0 m s^{-1}	desired velocity
τ	0.5 s	acceleration time
r	0.3 m	radius of pedestrians
A	$2 \times 10^3 \text{ N}$	interaction strength
B	0.08 m	interaction range
k	$1.2 \times 10^5 \text{ kg s}^{-2}$	bump effect
κ	$2.4 \times 10^9 \text{ kg m s}^{-1}$	friction effect

The interaction with the walls is treated analogously. By denoting d_{iW} as the distance to wall W , n_{iW} as the direction perpendicular to it, and t_{iW} as the direction tangential to it, we have

$$f_{iW} = \left\{ A_i \exp\left[\frac{(r_i - d_{iW})}{B_i}\right] + kg(r_i - d_{iW}) \right\} n_{iW} - \kappa g(r_i - d_{iW})(v_i \cdot t_{iW}) t_{iW}. \quad (4.5)$$

With the SFM described above, we simulate the collective escape of pedestrians towards a single exit in a room. Initially, a total of 100 pedestrians are randomly distributed within a square room measuring 10 m on each side. The room has an exit with a width of 2 m. The parameter values used to simulate the data are given in table 1, largely following [24]. We note here that in most of the real-world problems, the individuals are subject to different parameter values. To reduce the computational cost, we simplify the problem by assuming that the values for m_i , v_i^p , τ_i , r_i , A_i and B_i are identical for all individuals, as are in [10,11]. Among these parameters we assume that v^p , A and B are unknown and need to be inferred from the observation data. Namely, suppose we can observe the positions of pedestrians x_t at different times $t^{\text{obs}} = t \times \Delta_t$ for $t = 1, \dots, T$, and we aim to estimate the aforementioned parameters in the SFM with these sequential observations. The observation noise is assumed to be a zero-mean Gaussian with standard deviation σ . At every observation time, noise is added to the actual positions of each particle and the underlying distribution of this noisy observation is estimated by its empirical densities. We emphasize that we are able to observe the particle locations, but we are unable to track each particle. In this experiment, we take $\Delta_t = 0.1$ s and $T = 30$ and the simulation time step in SFM is $dt = 0.001$ s. The prior distribution for $\theta = (A, B, v^p)^T$ is taken to be uniform: $A \sim U[1200, 2200]$, $B \sim U[0.02, 0.2]$ and $v^p \sim U[0, 1.5]$. It is important to note that the SFM can encounter issues with unphysical parameter values that result in body compression and abrupt changes in velocities. Thus appropriate prior distributions of parameters are necessary to prevent such unrealistic simulations. Additionally, the simulation program includes an inspection procedure to identify and replace any unreasonable samples that may produce unrealistic signals during the simulation process. This approach enhances the stability of the simulation process and allows for more flexibility in setting parameter priors. We simulate the observation data with two different noise levels: $\sigma = 0.04$ and $\sigma = 0.1$, and use 500 samples in WD-SMCS to estimate parameters at each observation time. For each sample, the synthetic data $\hat{y}_t = \rho_t(x)$ is the approximated density of particle locations x_t at observation time t . These particle locations are generated by the SFM simulation with the parameters assigned the values of the corresponding samples. Importantly, in the simulation step, no noise is added to the particle locations due to our assumption that we have no prior knowledge about the measurement noise.

We first consider the small noise case where $\sigma = 0.04$. In figure 3, we plot the posterior means and the standard deviations against the number of iterations for all the three parameters. As can be seen from the figure, the mean values of all three parameters converge within 25 sequential observations, and the converged values are rather close to the ground truth, indicating good inference accuracy in the case when the observation noise is small. Recall that, a key element in the proposed method is that the simulated distribution should be close to the observed one. To facilitate the comparison of crowd distributions, we approximate the density map of particle locations using the kernel density estimation (KDE) [25] method. In figure 4, we compare the observed particle density and that




Figure 3. ($\sigma = 0.04$) Inference results of the SFM. Blue lines represent the average values of parameters estimated by 500 weighted samples and the standard deviations are shown as error bars. The horizontal lines represent the true parameter values.

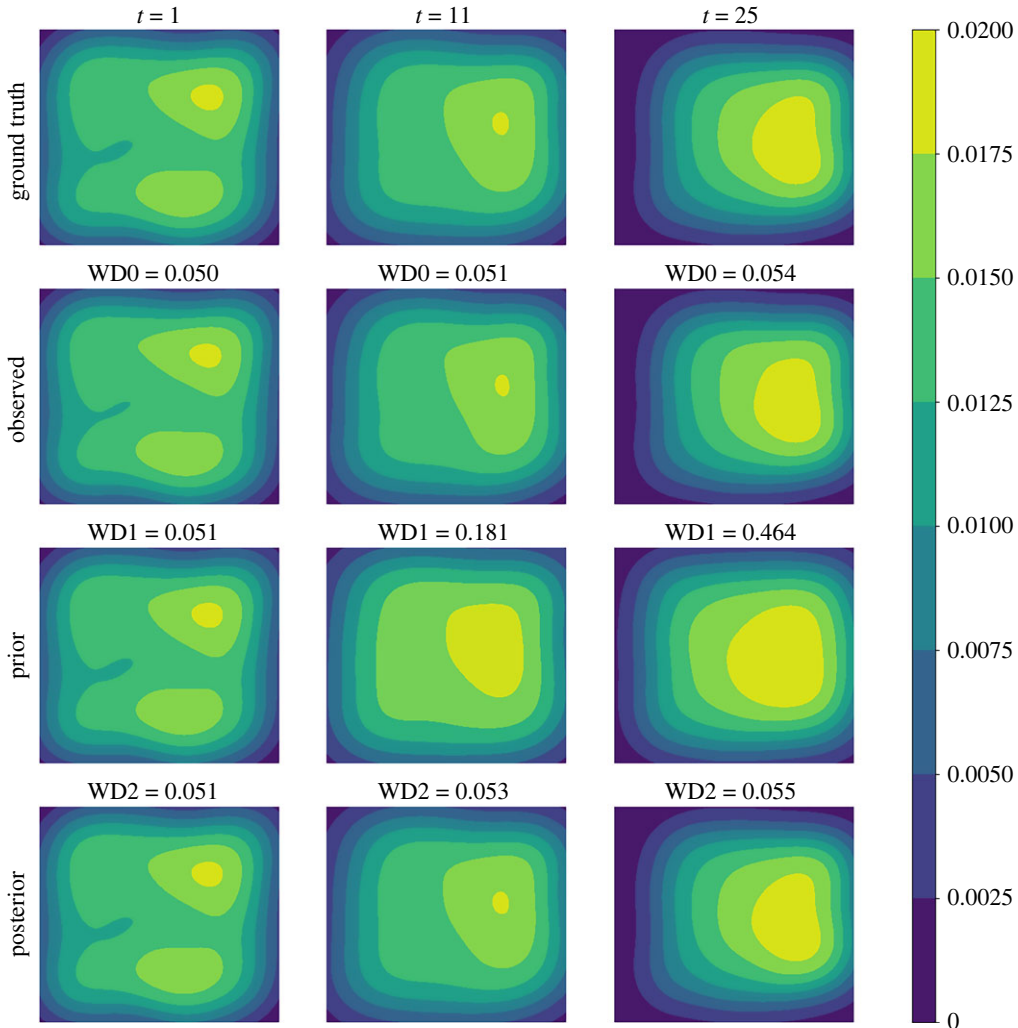


Figure 4. ($\sigma = 0.04$) Comparison among the ground truth (first row), the observation (second row), the prior prediction (third row) and the posterior prediction (last row) in terms of crowd density map at specific time points. The prior prediction is based on the mean values of parameter priors, while the posterior prediction is generated with the posterior means. WD0 is the WD between the ground truth and the observation. WD1 is the WD between the observed density and the prior-simulated one, and WD2 is that between the observed density and the posterior-generated one.

simulated with the posterior means (referred to as the posterior-simulated density) at three time steps: $t = 1, 11, 25$; we also show the densities from the prior means (referred to as the prior-simulated density) and the ground truth as a reference. For the two simulated densities, we also calculate the WD between them to the observed one. It can be seen that as the iteration proceeds (and therefore more data is available) the density associated with the prior means deviates from the observation, the posterior prediction density is closer to the observation with a smaller WD, compared to the prior prediction density. Next




Figure 5. ($\sigma = 0.1$) Inference results of the SFM. Blue lines represent the average values of parameters estimated by 500 weighted samples and the standard deviations are shown as error bars. The horizontal lines represent the true parameter values.




Figure 6. ($\sigma = 0.1$) Comparison among the ground truth (first row), the observation (second row), the prior prediction (third row) and the posterior prediction (last row) in terms of crowd density map at specific time points. The prior prediction is based on the mean values of parameter priors, while the posterior prediction is generated with the posterior means. WD0 is the WD between the ground truth and the observation. WD1 is the WD between the observed density and the prior-simulated one, and WD2 is that between the observed density and the posterior-generated one.

we consider the case of larger noise, i.e. $\sigma = 0.1$. As before we first plot the posterior means against the number of iterations in figure 5. One can see that, while the posterior means also approach to the true value as the iteration proceeds, more iterations (and hence more data) are required in order to accurately infer the parameters (especially A), due to the larger observation noise. We then show the same particle density plots as those in figure 4. One can see here that figure 6 is qualitatively similar to figure 4, which once again demonstrates that the proposed method can successfully obtain parameters values that can drive the particle density towards the observed one.




Figure 7. Inference results of the real-world data. Blue lines represent the average values of parameters estimated by 500 weighted samples and the standard deviations are shown as error bars.

We conduct further testing of the proposed algorithm using real-world data. Specifically, the experimental data [26] are collected in a corridor scenario leading straight to an entrance gate, which is common for concerts or other events. Initially, there are 75 participants waiting in the corridor, which has a width of 5.6 m. The exit is located along the negative y -axis and has a width of 0.5 m. Once instructed, participants start moving towards the entrance gate, and their locations are recorded during the process. More detailed information about the experiment can be found in [26]. We assume that the crowd dynamics in the corridor follow the SFM, and we use the observed distributions to infer the model's parameters. For this experiment, we take $\Delta t = 0.2$ s and $T = 40$. The parameters to be inferred and their priors are the same as before. In figure 7, we plot the posterior means and the standard deviations against the number of iterations for all the three parameters. From the figure, it can be observed that the mean value of the parameter v^p converges within 30 sequential observations, while the mean values of the other two parameters, A and B , remain highly variable. Given the unavailability of the actual model and parameter values, a reasonable practice is to assess whether the estimated model can capture the observed data patterns. Figure 8 shows the comparison between densities. One can see that both the prior-generated and the posterior-generated densities present significant discrepancies compared to the observed density. However, the posterior density demonstrates a smaller WD to the observed density compared to the prior density. We note that the performance is ultimately limited by SFM, which may not accurately describe this particular scenario. Furthermore, the simplification of assuming identical parameters for each participant also affects the performance. Nevertheless, we can still extract some useful information from the data using the proposed method.

4.2. The intelligent driver model

Our second example is the traffic on a four-lane highway (see figure 2, right), described by the IDM [27]. IDM is a widely adopted car-following model in microscopic traffic simulation. In this type of models, as its name suggests, the i -th vehicle follows the $(i-1)$ -th vehicle in front of it. The IDM assumes the acceleration of a follower i is a continuous function of a set of inputs—the follower's current velocity v_i ; the gap distance s_i to the leader; and the relative velocity (approach rate) Δv_i with respect to the preceding vehicle. The dynamics equation is defined as

$$\dot{v}_i = a^{(i)} \left[1 - \left(\frac{v_i}{v_0^{(i)}} \right)^\delta - \left(\frac{s^*(v_i, \Delta v_i)}{s_i} \right)^2 \right]. \quad (4.6)$$

in which

$$s_i = x_{i-1} - x_i - l^{(i)}, \quad (4.7)$$

$$\Delta v_i = v_i - v_{i-1} \quad (4.8)$$

and

$$s^*(v, \Delta v) = s_0^{(i)} + T_s^{(i)} v + \frac{v \Delta v}{2\sqrt{a^{(i)} b^{(i)}}}. \quad (4.9)$$

This expression can be understood as an interpolation of a free road acceleration $a_f = a^{(i)} [1 - (v_i/v_0^{(i)})^\delta]$ and an interaction deceleration $a_{int} = -a^{(i)} (s^*/s_i)^2$. Other constant parameters include the maximum desired velocity $v_0^{(i)}$, jam distance $s_0^{(i)}$, maximum desired acceleration $a^{(i)}$ and deceleration $b^{(i)}$, safe time headway $T_s^{(i)}$ and configurable exponent parameter δ , which controls the smoothness of the




Figure 8. Comparison among the observation (first row), the prior prediction (second row) and the posterior prediction (last row) in terms of crowd density map at specific time points. The prior prediction is based on the mean values of parameter priors, while the posterior prediction is generated with the posterior means. WD1 is the WD between the observed density and the prior-simulated one, and WD2 is that between the observed density and the posterior-generated one.

Table 2. List of parameter values in IDM and the three parameters to infer are in bold.

variable	value	description
$\mathbf{v_0}$	8.33 m s^{-1}	desired velocity
$\mathbf{T_s}$	1.6 s	safe time headway
\mathbf{a}	1.44 m s^{-1}	maximum acceleration
b	4.61 m s^{-1}	maximum deceleration
δ	4	acceleration exponent
s_0	2 m	jam distance
l	5 m	vehicle length

acceleration/deceleration. In this example, we study the case of identical vehicles whose model parameters $v_0^{(i)} = v_0$, $s_0^{(i)} = s_0$, $T_s^{(i)} = T_s$, $a^{(i)} = a$, $b^{(i)} = b$, $l^{(i)} = l$ are given in table 2.

With the IDM described above, we simulate the traffic flow on a four-lane highway. Each lane spans 300 m and starts empty. Vehicles arrive at a rate of three vehicles per second, which is intentionally chosen to ensure a sufficiently busy road, where interactions between vehicles occur frequently. Suppose that we are able to observe the vehicle locations x_t at a sequence of discrete-time points: $t^{\text{obs}} = t \times \Delta_t$ for $t = 1, \dots, T$, and we want to estimate the parameters $\theta = (v_0, a, T_s)^T$ in the IDM with these observed data with all other parameters known. The measurement noise is assumed to be a zero-mean Gaussian with standard deviation σ . This noise is added to the actual positions of each vehicle at every observation time. In this experiment, we take $\Delta_t = 1.0 \text{ s}$ and $T = 30$ and the simulation time step in IDM is $dt = 0.1 \text{ s}$. The prior distribution is taken to be uniform: $v_0 \sim U[5.56, 22.22]$, $a \sim U[0.5, 5]$ and $T_s \sim U[0.5, 4]$. We simulate the observation data with two different noise levels: $\sigma = 0.1$




Figure 9. ($\sigma = 0.1$) Inference results of the IDM. Blue lines show the posterior means of the parameters and the posterior standard deviations are shown as error bars. The horizontal lines represent the true parameter values.




Figure 10. ($\sigma = 1.0$) Inference results of the IDM. Blue lines show the posterior means of the parameters and the posterior standard deviations are shown as error bars. The horizontal lines represent the true parameter values.

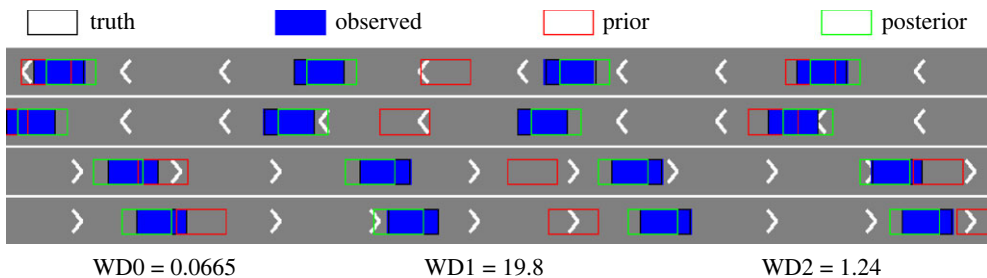


Figure 11. ($\sigma = 0.1$) The observed vehicle distribution (blue) compared against the ground truth (black), the posterior-simulated (green) and the prior-simulated (red) distributions over a specific interval on the highway. WDO is the WD between the observed distribution and the ground truth, WD1 is the WD between the observed distribution and the prior-simulated one, and WD2 is that between the observed distribution and the posterior-simulated one.

and $\sigma = 1.0$, and use 500 samples in WD-SMCS to estimate parameters at each observation time. For each sample, the synthetic data $\hat{y}_t = \rho_t(x)$ is the approximated density of vehicle locations x_t at observation time t . These vehicle locations are generated by the IDM simulation with the parameters assigned the values of the corresponding samples.

We first consider the small noise case with $\sigma = 0.1$ and plot the posterior means and standard deviations of the parameters against the number of iterations in figure 9. As can be seen from the figure, the mean estimated values of all three parameters converge to almost the ground truth within 15 sequential observations, which indicates the accurate inference in the case when the observation noise is small. The plots for the large noise case (i.e. $\sigma = 1.0$) are shown in figure 10, where we can see that it requires more data points to infer the parameters (especially a and T_s), but rather accurate results can still be obtained within 25 sequential observations. We also provide comparison of the vehicle distributions in figures 11 and 12 for the small and the large noise, respectively. In this example, comparing the distribution of vehicles based on their locations is more straightforward and visually interpretable due to the clear separation by lanes. Additionally, to accommodate the space limitations in the paper, we plot a subset of vehicle locations over a specific interval on the highway, covering approximately 100 m. Specifically, the figures show the observed vehicle distributions (blue), the posterior-simulated (green) and the prior-simulated (red) vehicle distributions. Visually in both figures the posterior-simulated distribution is rather close to the observation, while the prior-simulated one clearly deviates. Quantitatively the WD between the posterior-simulated distribution and the




Figure 12. ($\sigma = 1.0$) The observed vehicle distribution (blue) compared against the ground truth (black), the posterior-simulated (green) and the prior-simulated (red) distributions over a specific interval on the highway. WD0 is the WD between the observed distribution and the ground truth, WD1 is the WD between the observed distribution and the prior-simulated one, and WD2 is that between the observed distribution and the posterior-simulated one.

observed one is much smaller than that between the prior-simulated and the observed ones: they are 19.8 versus 1.24 in the small noise case and 30.2 versus 4.47 in the large noise case.

5. Conclusion

In summary, we consider in this work the problem of estimating parameters in a many-particle system. In particular, we aim to address the issue that in such problems it is often possible to observe the distributions of the particles rather than the trace of each individual particle. In this case, the likelihood function is not available and we adopt a likelihood-free SMCS method where the similarity (distance) between the observed data and the simulated data need to be characterized. Since the observed data is actually a distribution, we propose to use the WD to measure the similarity between the observed and the simulated data. Numerical experiments are also provided to demonstrate the performance of the proposed method. Our belief is that the proposed method has potential applications in a variety of real-world problems that involve the dynamics of multiple particles, particularly in the field of transport science. We intend to explore these potential applications in future research.

Data accessibility. No new data are presented in this study. Source code for the simulations and algorithms is available at <https://github.com/ChengChen0301/WD-SMCS>.

Authors' contributions. C.C.: conceptualization, data curation, formal analysis, methodology, validation, visualization, writing—original draft, writing—review and editing; L.W.: conceptualization, methodology, writing—original draft; J.L.: conceptualization, formal analysis, methodology, supervision, writing—original draft, writing—review and editing.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. We declare we have no competing interests.

Funding. No funding has been received for this article.

References

- Johansson A, Helbing D, Al-Abideen HZ, Al-Bosta S. 2008 From crowd dynamics to crowd safety: a video-based analysis. *Adv. Complex Syst.* **11**, 497–527. (doi:10.1142/S0219525908001854)
- Helbing D, Johansson A. 2011 In *Pedestrian, Crowd and Evacuation Dynamics*, pp. 697–716. New York, NY: Springer.
- Krauss S, Wagner P, Gawron C. 1997 Metastable states in a microscopic model of traffic flow. *Phys. Rev. E* **55**, 5597–5602. (doi:10.1103/PhysRevE.55.5597)
- Schranz M, Umlauf M, Sende M, Elmenreich W. 2020 Swarm robotic behaviors and current applications. *Front. Rob. AI* **7**, 36. (doi:10.3389/frobt.2020.00036)
- Johansson A, Helbing D, Shukla PK. 2007 Specification of the social force pedestrian model by evolutionary adjustment to video tracking data. *Adv. Complex Syst.* **10**, 271–288. (doi:10.1142/S0219525907001355)
- Hoogendoorn S, Hoogendoorn R. 2010 Calibration of microscopic traffic flow models using multiple data sources. *Phil. Trans. R. Soc. A* **368**, 4497–4517. (doi:10.1098/rsta.2010.0189)
- Paz A, Molano V, Martinez E, Gaviria C, Arteaga C. 2015 Calibration of traffic flow models using a memetic algorithm. *Transpor. Res. Part C: Emerg. Technol.* **55**, 432–443. (doi:10.1016/j.trc.2015.03.001)
- Gomes SN, Stuart AM, Wolfram MT. 2019 Parameter estimation for macroscopic pedestrian dynamics models from microscopic data. *SIAM J. Appl. Math.* **79**, 1475–1500. (doi:10.1137/18M1215980)
- Corbetta A, Muntean A, Vafayi K. 2015 Parameter estimation of social forces in pedestrian dynamics models via a probabilistic method. *Math. Biosci. Eng.* **12**, 337–356. (doi:10.3934/mbe.2015.12.337)
- Gödel M, Bode N, Köster G, Bungartz HJ. 2022 Bayesian inference methods to calibrate crowd dynamics models for safety applications. *Saf. Sci.* **147**, 105586. (doi:10.1016/j.ssci.2021.105586)
- Makinoshima F, Oishi Y. 2022 Crowd flow forecasting via agent-based simulations with sequential latent parameter estimation from aggregate observation. *Sci. Rep.* **12**, 11168. (doi:10.1038/s41598-022-14646-4)
- Csilléry K, Blum MG, Gaggiotti OE, François O. 2010 Approximate Bayesian computation (ABC) in practice. *Trends Ecol. Evol.* **25**, 410–418. (doi:10.1016/j.tree.2010.04.001)

13. Cranmer K, Brehmer J, Louppe G. 2020 The frontier of simulation-based inference. *Proc. Natl Acad. Sci. USA* **117**, 30055–30062. (doi:10.1073/pnas.1912789117)
14. Del Moral P, Doucet A, Jasra A. 2006 Sequential monte carlo samplers. *J. R. Stat. Soc. B (Stat. Methodol.)* **68**, 411–436. (doi:10.1111/j.1467-9868.2006.00553.x)
15. Toni T, Welch D, Strelkowa N, Ipsen A, Stumpf MP. 2009 Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *J. R. Soc. Interface* **6**, 187–202. (doi:10.1098/rsif.2008.0172)
16. Rubner Y, Tomasi C, Guibas L. 1998 A metric for distributions with applications to image databases. In *Sixth Int. Conf. on Computer Vision (IEEE Cat. No. 98CH36271)*, pp. 59–66.
17. Frogner C, Zhang C, Mobahi H, Araya-Polo M, Poggio T. 2015 Learning with a Wasserstein Loss. In *Proc. of the 28th Int. Conf. on Neural Information Processing Systems - Volume 2 NIPS'15*, pp. 2053–2061 Cambridge, MA. MIT Press.
18. Lee J, Raginsky M. 2018 Minimax statistical learning with wasserstein distances. In *Proc. of the 32nd Int. Conf. on Neural Information Processing Systems NIPS'18*, pp. 2692–2701. Red Hook, NY. Curran Associates Inc.
19. Bernton E, Jacob PE, Gerber M, Robert CP. 2019 Approximate Bayesian computation with the Wasserstein distance. *J. R. Stat. Soc. B (Stat. Methodol.)* **81**, 235–269. (doi:10.1111/rssb.12312)
20. Fearnhead P, Taylor BM. 2013 An adaptive sequential Monte Carlo sampler. *Bayesian Anal.* **8**, 411–438. (doi:10.1214/13-BA814)
21. Zhou Y, Johansen AM, Aston JA. 2016 Toward automatic model comparison: an adaptive sequential Monte Carlo approach. *J. Comput. Graph. Stat.* **25**, 701–726. (doi:10.1080/10618600.2015.1060885)
22. Robert C, Casella G. 2013 *Monte Carlo statistical methods*. New York, NY: Springer Science & Business Media.
23. Calvet L, Czellar V. 2011 State-observation sampling and the econometrics of learning models. *SSRN Electron. J.* (doi:10.2139/ssrn.1847646)
24. Helbing D, Farkas I, Vicsek T. 2000 Simulating dynamical features of escape panic. *Nature* **407**, 487–490. (doi:10.1038/35035023)
25. Rosenblatt M. 1956 Remarks on some nonparametric estimates of a density function. *Ann. Math. Stat.* **27**, 832–837. (doi:10.1214/aoms/1177728190)
26. Adrian J, Seyfried A, Sieben A. 2020 Crowds in front of bottlenecks at entrances from the perspective of physics and social psychology. *J. R. Soc. Interface* **17**, 20190871. (doi:10.1098/rsif.2019.0871)
27. Treiber M, Hennecke A, Helbing D. 2000 Congested traffic states in empirical observations and microscopic simulations. *Phys. Rev. E* **62**, 1805–1824. (doi:10.1103/PhysRevE.62.1805)