

Estimating ensemble flows on a hidden Markov chain

Isabel Haasler, Axel Ringh, Yongxin Chen, and Johan Karlsson

Abstract—We propose a new framework to estimate the evolution of an ensemble of indistinguishable agents on a hidden Markov chain using only aggregate output data. This work can be viewed as an extension of the recent developments in optimal mass transport and Schrödinger bridges to the finite state space hidden Markov chain setting. The flow of the ensemble is estimated by solving a maximum likelihood problem, which has a convex formulation at the infinite-particle limit, and we develop a fast numerical algorithm for it. We illustrate in two numerical examples how this framework can be used to track the flow of identical and indistinguishable dynamical systems.

I. INTRODUCTION

State tracking of a set of agents is an important issue in many areas, e.g., target tracking, see [4] and references therein. In this case, one is often interested in tracking one single or a set of multiple distinct targets. However, in many applications information for each agent may not be available, e.g., if the population is too large to track every single agent, as in many biological systems, or due to data privacy [24]. In this work, we thus consider tracking the evolution of a finite ensemble of indistinguishable agents. Based on reduced and incomplete measurements of the whole population at different time points, we aim to recover an estimate of the discrete-time flow of the ensemble. Related state estimation problems for a continuum of agents and in continuous time have been considered in [10], [35] (see also [6]).

In this work, we use a hidden Markov model (HMM) to describe the particle flows and aggregate observations, similar to [2], and seek the most likely paths that the agents have taken. These paths are found by maximizing the log-likelihood function of the flow, subject to the constraint that the flow matches the given measurements. This gives rise to a convex maximum entropy type optimization problem, and we derive an efficient algorithm for solving it.

The problem of finding the most likely path for the evolution of a distribution is related to a discrete Schrödinger bridge problem [28].

Schrödinger's thought experiment [31] has indirectly given rise to the concept of reciprocal processes [5], [22], [25], which connects this work to tracking of moving objects using reciprocal processes [18], [32], [34]. However, as mentioned before, we consider estimating the flow of an ensemble rather than single target tracking.

This work was supported by the Swedish Research Council (VR), grant 2014-5870, SJTU-KTH cooperation grant and the NSF under grant 1901599.

I. Haasler, A. Ringh, and J. Karlsson are with the Division of Optimization and Systems Theory, Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden. {haasler, aringh}@kth.se, johan.karlsson@math.kth.se

Y. Chen is with the School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, USA. yongchen@gatech.edu

The outline of the paper is as follows: Section II presents background material, in particular on HMMs, Schrödinger bridges, and optimal mass transport. In section III we derive the maximum likelihood problem for a Markov chain with a known initial and final distribution, and relate it to prior work on the Schrödinger bridge problem [28]. In Section IV, which contains the main contribution, we extend this maximum entropy framework to HMMs with indirect and noisy observations. Moreover, we derive the corresponding maximum likelihood problem, and develop a fast iterative algorithm for solving it. The method is demonstrated on two examples in section V, and section VI contains conclusions and future directions. Some proofs are deferred to the appendix for improved readability.

II. BACKGROUND

A. Notation

By \cdot / \cdot , \odot , $\log(\cdot)$, and $\exp(\cdot)$ we denote elementwise division, multiplication, logarithm, and exponential of matrices and vectors. Moreover, by $\text{supp}(\cdot)$ we denote the support of a matrix, i.e., the non-zero elements.

B. Hidden Markov chains

In this work, we consider hidden Markov models for stochastic modeling of a group of indistinguishable agents/particles. For an introduction to HMMs, see, e.g., [21], [29]. An HMM is a structure that consists of two stochastic processes. The first part is a Markov chain that evolves over a hidden set of states $X = \{X_1, X_2, \dots, X_n\}$ and is used to model the unobserved, underlying state of the system. We denote the state at time t by q_t . The stochastic state transitions are encoded in the state transition matrix $A = [a_{ij}]_{i,j=1}^n$, where $a_{ij} = P(q_{t+1} = X_j | q_t = X_i)$. The second part is an observation process providing partial and noisy information of the underlying process; here we use the observation symbols $Y = \{Y_1, Y_2, \dots, Y_m\}$. Moreover, the observation process is also Markovian with respect to the underlying state in the hidden Markov chain, i.e., the observation probabilities can be summarized in a matrix $B \in \mathbb{R}^{n \times m}$ with elements $b_{jk} = P(Y_k \text{ at } t | q_t = X_j)$.

C. Schrödinger bridges and large deviations

In the early 1930s, Schrödinger discussed the problem of determining the evolution of particles between two observed distributions [31]. Assuming a cloud of independent Brownian particles is observed at time instance $t = 0$, the expected distribution at $t = 1$ would be described by

$$\rho_1(x_1) = \int_{\mathbb{R}^n} q_\epsilon(0, x_0, 1, x_1) \rho_0(x_0) dx_0, \quad (1)$$

where q_ϵ is the Brownian transition probability kernel

$$q_\epsilon(s, x, t, y) = \frac{1}{(2\pi(t-s)\epsilon)^{n/2}} \exp\left(-\frac{\|x-y\|^2}{2(t-s)\epsilon}\right),$$

and where the parameter ϵ denotes a diffusion coefficient. Schrödinger studied the problem where the observed particle distribution differs from the expected distribution (1). The most likely particle evolution connecting, hence bridging, the two marginals is called the Schrödinger bridge.

The Schrödinger bridge problem was later formulated in the context of large deviation theory [19, Sec. II.1.3], the study of rare events in the sense of deviations from the law of large numbers [15], [16]. As the number of trials (or particles) goes to infinity, the probability of such rare events approaches zero. Large deviation theory studies the rate of this decay, which can often be characterized by the exponential of a so called rate function.

Modeling the particle evolutions as independent identically distributed random variables on path space, a Schrödinger bridge is a probability measure \mathcal{P} on path space that is most likely to describe the rare event of observing the two particle distributions. Such a measure is obtained by minimizing the corresponding rate function, which turns out to be the relative entropy with respect to the underlying probability law of the Brownian motion. In other words, \mathcal{P} is the measure that is “most similar” to the Wiener measure \mathcal{W} in the sense that it minimizes the relative entropy [26]

$$H(\mathcal{P} \mid \mathcal{W}) = \int \log\left(\frac{d\mathcal{P}}{d\mathcal{W}}\right) d\mathcal{P} \quad (2)$$

over all probability measures that are absolutely continuous with respect to \mathcal{W} and have the given particle distributions as marginals. The Schrödinger bridge can be constructed from the solution to a certain system of equations, called the Schrödinger system. A space and time discrete Schrödinger bridge problem for Markov chains is analysed in [8], [20], [28].

D. Optimal mass transport

Another recently established connection of Schrödinger bridges is to the problem of optimal mass transport (OMT) [7], [9], [26], [27]. As the diffusion coefficient ϵ in (1) approaches 0, the solution to the Schrödinger bridge problem tends to the solution to a corresponding optimal mass transport problem [26]. Moreover, the Schrödinger bridge formulation is a regularization of OMT, as it is strictly convex and therefore guarantees a unique solution.

We introduce a discretized formulation of the OMT problem. For an extensive discussion of OMT see, e.g., [33]. Consider a discretization $\{x_1, \dots, x_n\}$ of a compact space X and two distributions $\mu_0, \mu_1 \in \mathbb{R}^n$ defined on this discretization. Given a cost matrix $C = [c_{ij}]_{i,j=1}^n$, where c_{ij} denotes the cost of transporting a unit mass from point x_i to x_j , we seek a transport plan $M = [m_{ij}]_{i,j=1}^n$, where m_{ij} denotes the amount of mass being transported from x_i to x_j , that minimizes the total transportation cost $\text{tr}(C^T M)$ between the two distributions, i.e., the transport plan is

required to satisfy $M\mathbf{1} = \mu_0$ and $M^T\mathbf{1} = \mu_1$, where $\mathbf{1}$ denotes an $n \times 1$ -vector of ones.

Solving this linear program is computationally expensive for large n . It was therefore proposed to regularize the problem by introducing a Kullback-Leibler divergence term (sometimes called entropy term) to the objective [12].

Definition 1: Let p and q be two nonnegative vectors or matrices of the same dimension. The Kullback-Leibler (KL) divergence between p from q is defined as

$$H(p|q) := \sum_i p_i \log\left(\frac{p_i}{q_i}\right)$$

where $0 \log 0$ is defined to be 0. Note that $H(p|q)$ is jointly convex over p, q . See, e.g., [11] for more properties and interpretation of the KL divergence.

The discretized and regularized OMT problem then reads

$$\begin{aligned} & \underset{M \in \mathbb{R}^{n \times n}}{\text{minimize}} \quad \text{trace}(C^T M) + \epsilon H(M \mid \mathbf{1}_{n \times n}) \\ & \text{subject to} \quad M\mathbf{1} = \mu_0, \quad M^T\mathbf{1} = \mu_1, \end{aligned} \quad (3)$$

where $\epsilon > 0$ is a regularization parameter and $\mathbf{1}_{n \times n}$ denotes an $n \times n$ -matrix of ones. The solution to this problem may be found by Sinkhorn iterations, which correspond to the fixed point iteration for the Schrödinger system in [7].

III. PARTICLE DYNAMICS OVER A MARKOV CHAIN

Consider a cloud of N particles, where each particle is evolving according to a Markov chain as described in Section II-B. Let the vectors $\mu_t \in \mathbb{N}^n$ describe the particle distributions at time $t \in \{0, 1\}$, where the i -th element $(\mu_t)_i$ denotes the number of particles in state X_i at time t . In analogy to the OMT framework, we define the mass transfer matrix $M = [m_{ij}]_{i,j=1}^n$, where m_{ij} denotes the number of particles that transit from state X_i to state X_j . Note that the mass transport matrix satisfies $M\mathbf{1} = \mu_0$ and $M^T\mathbf{1} = \mu_1$.

The state transition matrix $A = [a_{ij}]_{i,j=1}^n$ contains the particle transition probabilities. Thus, given the initial state μ_0 , the probability of a mass transfer matrix M is

$$P_{\mu_0, A}(M) = \prod_{i=1}^n \left(\binom{(\mu_0)_i}{m_{i1}, m_{i2}, \dots, m_{in}} \prod_{j=1}^n a_{ij}^{m_{ij}} \right), \quad (4)$$

where $\binom{\cdot}{\cdot, \dots, \cdot}$ denotes a multinomial coefficient. The expected distribution at time $t = 1$ is then given by $E(\mu_1 | \mu_0) = A^T \mu_0$. If μ_1 is observed to be different from $A^T \mu_0$ a discrete version of the Schrödinger bridge problem can be solved (see Section II-C). That is, to find the matrix M that maximizes $P_{\mu_0, A}(M)$ subject to that the constraint $M^T\mathbf{1} = \mu_1$ on the final marginal is satisfied.

If the number of particles is large, then the log-likelihood of (4) can be approximated in terms of a KL divergence.

Proposition 1: Given A , let $\mu_0^{(N)} \in \mathbb{N}^n$ be a sequence of distributions with N particles, and $M^{(N)} \in \mathbb{N}^{n \times n}$ be a sequence of mass transfer matrices such that $M^{(N)}\mathbf{1} = \mu_0^{(N)}$ and $\text{supp}(M^{(N)}) \subseteq \text{supp}(\text{diag}(\mu_0^{(N)})A)$. Then there exists a

constant $C > 0$ such that for all N it holds that

$$\left| \log \left(P_{\mu_0^{(N)}, A}(M^{(N)}) \right) + H(M^{(N)} | \text{diag}(\mu_0^{(N)})A) \right| \leq C \log(N).$$

Proof: See appendix. \blacksquare

Proposition 1 implies that for sequences $\mu_0^{(N)}$ and $M^{(N)}$ satisfying the assumptions, if

$$\frac{1}{N} \mu_0^{(N)} \rightarrow \bar{\mu}_0 \quad \text{and} \quad \frac{1}{N} M^{(N)} \rightarrow \bar{M}$$

as $N \rightarrow \infty$, then

$$\frac{1}{N} \log \left(P_{\mu_0^{(N)}, A}(M^{(N)}) \right) \rightarrow -H(\bar{M} | \text{diag}(\bar{\mu}_0)A)$$

as $N \rightarrow \infty$. This means that the KL divergence approximates the log-likelihood of $P_{\mu_0, A}(M)$ with increasing accuracy as the number of particles increases. We write this informally as

$$P_{\mu_0, A}(M) \sim e^{-H(M | \text{diag}(\mu_0)A)}.$$

In terms of large deviation theory, we thus interpret $H(\cdot | \text{diag}(\mu_0)A)$ as the rate function for $P_{\mu_0, A}(\cdot)$. In fact, Proposition 1 can also be derived from a large deviation principle (see, e.g., [15, Ch. 2.1.1]).

For systems with many particles, we may therefore formulate the problem of finding the most likely mass transfer matrix M between distributions μ_0 and μ_1 with underlying state transition matrix A as the convex optimization problem

$$\begin{aligned} & \underset{M \in \mathbb{R}^{n \times n}}{\text{minimize}} \quad H(M | \text{diag}(\mu_0)A) \\ & \text{subject to} \quad M\mathbf{1} = \mu_0, \quad M^T\mathbf{1} = \mu_1. \end{aligned} \quad (5)$$

Remark 1: Let A and μ_0 be strictly positive. With the cost matrix

$$C = -\epsilon \log(\text{diag}(\mu_0)A),$$

the entropy regularized OMT problem (3) is equivalent to problem (5). Note that entropy regularized OMT problems have previously been solved by formulating them in terms of KL-projection problems [1].

A. Connection to Schrödinger bridges

We note that given the prior distribution μ_0 , the objective in (5) may be written as

$$H(M | A) - H(\mu_0 | \mathbf{1})$$

where the second term is constant. Hence, if we associate A and M with the measures $d\mathcal{W}$ and $d\mathcal{P}$ in (2), the problem in Proposition 1 relates to a discretized Schrödinger bridge. Our problem formulation indeed corresponds to the time and space discrete Schrödinger bridge from [28]. To see this, consider a Markov chain of length T . Using Proposition 1, knowing the marginals μ_0 and μ_T , we can find the most likely evolution of the particles between them as the solution to

$$\begin{aligned} & \underset{M_{[1:T]}, \mu_{[1:T-1]}}{\text{minimize}} \quad \sum_{t=1}^T H(M_t | \text{diag}(\mu_{t-1})A) \\ & \text{subject to} \quad M_t\mathbf{1} = \mu_{t-1}, \quad M_t^T\mathbf{1} = \mu_t, \\ & \quad \text{for } t = 1, \dots, T. \end{aligned} \quad (6)$$

Note that for a nonnegative matrix M_t and strictly positive marginal μ_{t-1} , the first constraint asserts that there is a row-stochastic matrix \bar{M}_t such that $M_t = \text{diag}(\mu_{t-1})\bar{M}_t$. Plugging this expression for the matrices M_t into (6) gives

$$\begin{aligned} & \underset{\bar{M}_{[1:T]}, \mu_{[1:T-1]}}{\text{minimize}} \quad \sum_{t=1}^T \sum_i (\mu_{t-1})_i H((\bar{M}_t)_{i \cdot}, A_{i \cdot}) \\ & \text{subject to} \quad \bar{M}_t\mathbf{1} = \mathbf{1}, \quad \mu_t = \bar{M}_t^T \mu_{t-1}, \\ & \quad \text{for } t = 1, \dots, T. \end{aligned} \quad (7)$$

Here $A_{i \cdot}$ denotes the i -th row of A . This is precisely the formulation of a Schrödinger bridge over a Markov chain from [28, eq. (24)] with time invariant transition probabilities. In [28] it is shown that a unique solution to a corresponding Schrödinger system exists if μ_T is a strictly positive distribution and all elements are strictly positive in the matrix A raised to the power T . The solution to the Schrödinger system may be obtained by a fixed point iteration [20], which is linked to the Sinkhorn iterations for entropy regularized OMT problems, cf. Section II-D.

We note that the optimization problem (7) is non-convex and will thus work with the formulation (6) in the remaining part of this article.

IV. PARTICLE DYNAMICS OVER HIDDEN MARKOV CHAIN

In this section, we extend our framework to the setting of a hidden Markov chain. The initial marginal distribution is assumed to be known. In case the hidden states are linked to the observations by a deterministic linear mapping, they may be estimated in a similar fashion as in [17]. Here instead, we consider the non-deterministic case where the available observations emerge from the hidden distributions through an observation probability matrix $B \in \mathbb{R}^{n \times m}$.

Equivalently to the mass transfer plans M , define the observation matrix $D \in \mathbb{N}^{n \times m}$, with entries d_{jk} denoting the number of particles that are in hidden state X_j and observed in state Y_k . Given a hidden state μ , the probability for any observation matrix D is given by $P_{\mu, B}(D)$ as defined in (4). Hence, the large deviation result in Proposition 1 holds for D with rate function $H(\cdot | \text{diag}(\mu)B)$.

Given an initial distribution $\mu_0 \in \mathbb{N}^n$ and a set of measurements $\Phi_1, \dots, \Phi_T \in \mathbb{N}^m$, we seek the most likely set of matrices M_1, \dots, M_T and D_1, \dots, D_T such that for some set of hidden distributions μ_1, \dots, μ_T it holds that

$$\begin{aligned} M_t\mathbf{1} &= \mu_{t-1}, \quad M_t^T\mathbf{1} = \mu_t, \\ D_t\mathbf{1} &= \mu_t, \quad D_t^T\mathbf{1} = \Phi_t, \quad \text{for } t = 1, \dots, T. \end{aligned} \quad (8)$$

This model is illustrated in Figure 1. The maximum likelihood solution is obtained by solving the optimization problem

$$\underset{M_{[1:T]}, D_{[1:T]}, \mu_{[1:T]}}{\text{maximize}} \quad \prod_{t=1}^T P_{\mu_{t-1}, A}(M_t) P_{\mu_t, B}(D_t)$$

subject to (8). From Proposition 1 it follows that

$$\log \left(\prod_{t=1}^T P_{\mu_{t-1}, A}(M_t) P_{\mu_t, B}(D_t) \right)$$

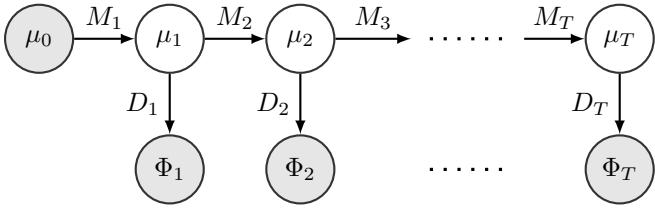


Fig. 1: Illustration of the hidden Markov model corresponding to (8).

can be approximated by

$$-\sum_{t=1}^T \left(H(M_t | \text{diag}(\mu_{t-1})A) + H(D_t | \text{diag}(\mu_t)B) \right) \quad (9)$$

when the number of particles is large. We thus estimate the matrices M_1, \dots, M_T , D_1, \dots, D_T and the hidden states μ_1, \dots, μ_T by maximizing (9) subject to the constraints (8), i.e., by solving

$$\begin{aligned} & \underset{M_{[1:T]}, D_{[1:T]}, \mu_{[1:T]}}{\text{minimize}} \quad \sum_{t=1}^T H(M_t | \text{diag}(\mu_{t-1})A) \\ & \quad + \sum_{t=1}^T H(D_t | \text{diag}(\mu_t)B) \quad (10) \\ & \text{subject to} \quad M_t \mathbf{1} = \mu_{t-1}, \quad M_t^T \mathbf{1} = \mu_t \\ & \quad D_t \mathbf{1} = \mu_t, \quad D_t^T \mathbf{1} = \Phi_t \\ & \quad \text{for } t = 1, \dots, T. \end{aligned}$$

Remark 2: The modeling assumptions leading to this optimization problem require knowledge of the initial distribution as well as the transition and observation probabilities for the hidden Markov model. However, in practice these are typically not known exactly. In the examples in Section V, we illustrate that the estimation is accurate even when there are significant model errors. The generalization of the proposed method to the case where the initial distribution is not known will be discussed in a forthcoming paper.

A. Computational method

In this section we develop a numerical method to solve problem (10). To this end, recall from Remark 1 the connection between KL-minimization problems and entropy regularized OMT problems. The latter can be efficiently solved by Sinkhorn iterations [12], which in turn are equivalent to a block coordinate ascent in a dual problem [23]. Motivated by this, we choose to follow a similar approach.

Proposition 2: Let $u_1 \in \mathbb{R}^n$ and $v_t \in \mathbb{R}^m$, for $t = 1, \dots, T$, be positive initial values, and iterate the following steps:

- (1) $u_1 = e\mathbf{1}./(Aw_1)$
- (2) $v_t = e\Phi_t./(B^T(y_t \odot (Aw_{t+1})))$ for $t = 1, \dots, T$,

where in each step and for each t in step (2), the vectors y_t and w_t are recursively defined as

$$\begin{aligned} y_1 &= A^T(\mu_0 \odot u_1), \\ y_t &= A^T(y_{t-1} \odot (Bv_{t-1})), \quad t = 2, \dots, T \end{aligned}$$

and

$$\begin{aligned} w_T &= Bv_T \\ w_t &= (Bv_t) \odot (Aw_{t+1}), \quad t = 1, \dots, T-1. \end{aligned}$$

In the limit point of the iteration, the estimates for the hidden marginals are then recursively constructed, starting from the known μ_0 , as

$$\mu_t = \text{diag}(w_t) A^T (\mu_{t-1} ./ (Aw_t)), \quad t = 1, \dots, T.$$

Furthermore, the corresponding mass transfer matrices are given by

$$\begin{aligned} M_t &= \frac{1}{e} \text{diag}(\mu_{t-1} \odot u_t) A \text{diag}(w_t), \\ D_t &= \frac{1}{e} \text{diag}(\mu_t \odot x_t) B \text{diag}(v_t), \end{aligned}$$

where

$$\begin{aligned} x_t &= e\mathbf{1} ./ (Bv_t) \\ u_t &= e\mathbf{1} ./ (Aw_t) \end{aligned}$$

for $t = 1, \dots, T$.

Proof: See appendix. ■

It is worth noting that intermediate results of the vectors y_t and w_t may be stored, such that the update of u_1 requires only one matrix-vector multiplications with A , and the update of v_t , for any $t = 1, \dots, T$, involves two multiplications with B and one with A . One iteration sweep, i.e. one update of u_1 and the set v_t , for $t = 1, \dots, T$, thus requires $\mathcal{O}(Tn \max(n, m))$ operations.

V. SIMULATIONS

A. Particle dynamics

Consider a cloud of 1000 particles evolving from an initial distribution $\mu_0 \in \mathbb{R}^n$ with $n = 100$ states. The particles transition matrix is given by $\tilde{A} \in \mathbb{R}^{n \times n}$ with elements

$$\tilde{a}_{ij} \sim \exp \left(\frac{1}{2\sigma_{\tilde{a}}^2} (i - j - 1)^2 \right), \quad \text{with } \sigma_{\tilde{a}} = 0.5,$$

which corresponds to a discretization of a normal distribution $\mathcal{N}(1, 0.5)$, and thus induces a drift on the dynamics of the cloud. The true dynamics of the particles are assumed to be unknown and instead modeled by a transition matrix $A \in \mathbb{R}^{n \times n}$ with elements

$$a_{ij} \sim \exp \left(\frac{1}{2\sigma_a^2} (i - j)^2 \right), \quad \text{with } \sigma_a = 2.$$

At each time instance, the particles are observed in $m = 5$ bins, where the observation probability matrix $B \in \mathbb{R}^{n \times m}$ has elements

$$b_{ij} \sim \exp \left(\frac{1}{2\sigma_b^2} \left(j - \frac{i + 10}{20} \right)^2 \right), \quad \text{with } \sigma_b = 0.5.$$

We estimate the flow of the particles and hidden particle distributions solving problem (10) for $T = 50$ time instances with the method proposed in Proposition 2. One estimate is formed using the true initial distribution μ_0 as a prior distribution, and for a second estimate we use a uniform prior.

Figure 2 shows the true hidden particle cloud, the corresponding observations, and the two estimates. With full information of the initial states available, the proposed method provides a good estimate of the hidden states despite discrepancies between the true and assumed transition matrices \tilde{A} and A . In the case of no prior information, i.e., the prior distribution is set to be uniform, we see that the estimate converges to the estimate with true prior within a few time steps. This indicates that the proposed method is robust to modeling uncertainties and lack of information in the initial state.

B. Tracking ensembles over a network

In this example we consider the problem of tracking a number of indistinguishable agents over a network given measurements from sensors that are distributed around the network. This is inspired by [13], where an HMM is used to estimate the flow of a crowd in an urban environment based on observations generated when cell phones connect to Wi-Fi sensors. The environment is modeled as a network of nodes and arcs, where the arcs represent walking paths in the area and the nodes are the intersections between the paths.

For this application, the optimization problem (10) needs to be extended to allow for multiple measurements. To this end, let Φ_{st} be a set of observations from measurement unit s at time point t , for $t = 1, \dots, T$, and for $s = 1, \dots, S$. We obtain the maximum likelihood solution as the optimal solution to¹

$$\begin{aligned} & \underset{M_{[1:T]}, D_{[1:T], [1:S]}, \mu_{[1:T]}}{\text{minimize}} \quad \sum_{t=1}^T H(M_t \mid \text{diag}(\mu_{t-1}) A) \\ & \quad + \sum_{t=1}^T \sum_{s=1}^S H(D_{st} \mid \text{diag}(\mu_t) B_s) \\ \text{subject to} \quad & M_t \mathbf{1} = \mu_{t-1}, \quad M_t^T \mathbf{1} = \mu_t \quad (11) \\ & D_{st} \mathbf{1} = \mu_t, \quad D_{st}^T \mathbf{1} = \Phi_{st} \\ & \text{for } t = 1, \dots, T, \text{ and } s = 1, \dots, S. \end{aligned}$$

Consider a hidden Markov model where the states $X = \{X_1, \dots, X_n\}$ are the edges in the directed graph $\mathcal{G} = (V, X)$, and where the edge $X_i = (V_i^{\text{in}}, V_i^{\text{out}})$ goes from $V_i^{\text{in}} \in V$ to $V_i^{\text{out}} \in V$. In this example we will use the graph illustrated in Figure 3, consisting of 11 nodes and $n = 28$ edges. For the true model, transition probabilities are defined according to weights in the graph that represent which walking paths are preferred by the pedestrians. For the model used in the estimation we assume that this information is not known and use uniform weights. More specifically, the

¹This is a convex optimization problem which in principle can be solved with off-the-shelf solvers. In this example we use an efficient algorithm in the spirit of Proposition 2, but due to lack of space we defer the exact algorithm to a forthcoming paper.

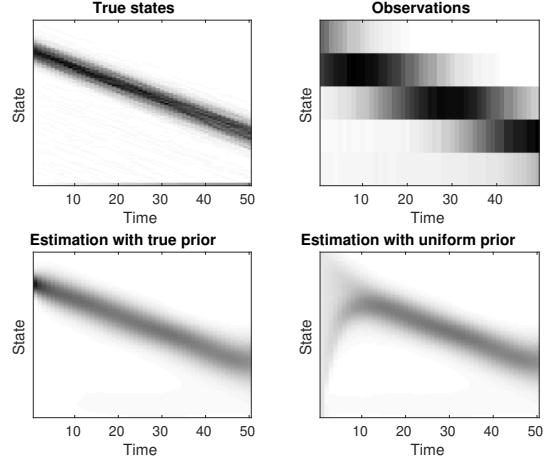


Fig. 2: Particle cloud reconstructed from observations.

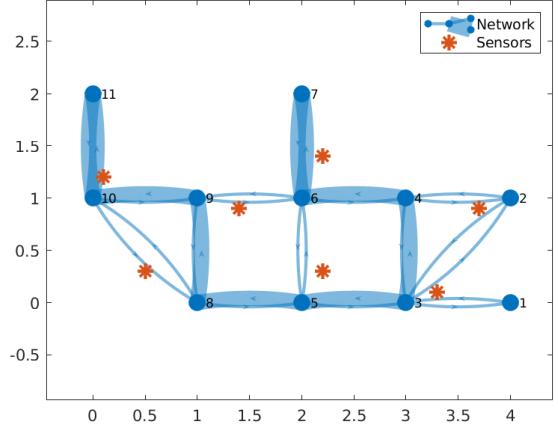


Fig. 3: Network and sensors.

transition probabilities are given by

$$\tilde{a}_{ij} = \begin{cases} 0.5, & \text{if } j = i \\ 0.5 w_{ij} \left(\sum_{\{k: V_k^{\text{in}} = V_i^{\text{out}}\}} w_{ik} \right)^{-1} & \text{if } V_j^{\text{in}} = V_i^{\text{out}} \\ 0, & \text{else,} \end{cases}$$

where $\{w_{ij}\}$ is a set of weights. For the transition matrix A used in the estimation we assign uniform weights $w_{ij} = 1$, for all (i, j) with $V_j^{\text{in}} = V_i^{\text{out}}$. For the true transitions the weights are defined as

$$w_{ij} = \begin{cases} 20, & X_j \in \mathcal{W} \text{ and } V_j^{\text{out}} \neq V_i^{\text{in}} \\ 0, & V_j^{\text{out}} = V_i^{\text{in}} \\ 1, & \text{else,} \end{cases}$$

for (i, j) such that $V_j^{\text{in}} = V_i^{\text{out}}$, and where \mathcal{W} is the set of edges highlighted in Figure 3. Note that the second case implies that agents do not transition to the reverse edge in the next time step.

The agents are observed by a set of $S = 7$ sensors located at the positions indicated in Figure 3. The observation

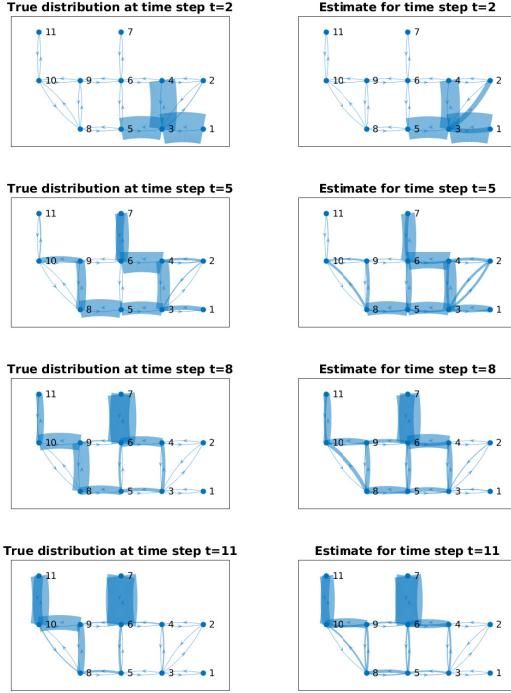


Fig. 4: Ensemble flow over network.

probability for an agent on edge X_i to be detected by a given sensor s is defined as $b_{i1}^s = \min(0.99, 2e^{-5d})$, where d denotes the Euclidean distance between the location of s and the midpoint of X_i . Consequently the probability of not being detected is $b_{i2}^s = 1 - b_{i1}^s$.

Given an initial distribution of 100 agents on the edge $(1, 3)$, the flow and the measurements for the true ensemble are computed for $T = 20$ time steps using the true transition and observation probabilities. Then we estimate the flow by solving the optimization problem (11). The true and estimated particle distributions are compared for some time steps in Figure 4, where the width of each edge is proportional to the number of agents on it. As can be seen in the figure the proposed method provides a good estimate also for this example.

VI. CONCLUSIONS AND FUTURE DIRECTIONS

In this work, we propose a method for estimating the flow of an ensemble of particles on a hidden Markov chain. The estimation, which is formulated as a maximum likelihood problem, can be recast as a convex optimization problem, for which we provide an efficient algorithm.

There are several natural directions in which this work can be extended. One restriction in this paper is that the number of agents are fixed and known, thus extensions with a variable number of particles, such as birth/death processes, could be of interest (cf. [14]). Furthermore, the model may be extended to continuous-state dynamics. Another natural direction is to study the connections to reciprocal processes.

A. Proof of Proposition 1

Let $\mu_0^{(N)}$ and $M^{(N)}$ be as described in the statement of the proposition. Moreover, let $\mathcal{Z}^{(N)} = \{(i, j) \mid m_{ij}^{(N)} \neq 0\}$, which is non-empty, and let $\mathcal{Z}_i^{(N)} = \{j \mid (i, j) \in \mathcal{Z}^{(N)}\}$, which is non-empty if and only if $(\mu_0^{(N)})_i > 0$. Furthermore, let $\mathcal{Y}^{(N)} = \{i \mid (\mu_0^{(N)})_i > 0\}$. Then, using Stirling's formula

$$\sqrt{2\pi} n^{n-1/2} e^{-n} \leq n! \leq e n^{n-1/2} e^{-n},$$

see, e.g., [30], for $i \in \mathcal{Y}^{(N)}$ the multinomial coefficient in (4) can be bounded from above by

$$\begin{aligned} & \binom{(\mu_0^{(N)})_i}{m_{i1}^{(N)}, m_{i2}^{(N)}, \dots, m_{in}^{(N)}} \\ & \leq e^{-(|\mathcal{Z}_i^{(N)}|-1)} \exp \left(\sum_{j \in \mathcal{Z}_i^{(N)}} m_{ij}^{(N)} - (\mu_0^{(N)})_i \right) \\ & \quad \cdot (\mu_0^{(N)})_i^{(\mu_0^{(N)})_i + \frac{1}{2}} \prod_{j \in \mathcal{Z}_i^{(N)}} (m_{ij}^{(N)})^{-(m_{ij}^{(N)} + \frac{1}{2})} \\ & \leq (\mu_0^{(N)})_i^{(\mu_0^{(N)})_i + \frac{1}{2}} \prod_{j \in \mathcal{Z}_i^{(N)}} (m_{ij}^{(N)})^{-(m_{ij}^{(N)} + \frac{1}{2})}, \end{aligned}$$

where $|\mathcal{Z}_i^{(N)}|$ denotes the cardinality of the set, and where the second inequality follows from the fact that $\sum_{j \in \mathcal{Z}_i^{(N)}} m_{ij} = \sum_{j=1}^n m_{ij} = (\mu_0)_i$, and that $e^{-(|\mathcal{Z}_i^{(N)}|-1)} \leq 1$. Thus, the log-likelihood of the probability for a transfer plan $M^{(N)}$ can be upper-bounded as follows:

$$\begin{aligned} & \log \left(P_{\mu_0^{(N)}, A} (M^{(N)}) \right) \\ & \leq \sum_{(i, j) \in \mathcal{Z}^{(N)}} \left(m_{ij}^{(N)} \log(a_{ij}) - \left(m_{ij}^{(N)} + \frac{1}{2} \right) \log(m_{ij}^{(N)}) \right) \\ & \quad + \sum_{i \in \mathcal{Y}^{(N)}} \left((\mu_0^{(N)})_i + \frac{1}{2} \right) \log((\mu_0^{(N)})_i) \\ & = \sum_{(i, j) \in \mathcal{Z}^{(N)}} \left(m_{ij}^{(N)} \log \left(\frac{(\mu_0^{(N)})_i a_{ij}}{m_{ij}^{(N)}} \right) \right) \\ & \quad - \sum_{(i, j) \in \mathcal{Z}^{(N)}} \frac{1}{2} \log(m_{ij}^{(N)}) + \sum_{i \in \mathcal{Y}^{(N)}} \frac{1}{2} \log((\mu_0^{(N)})_i) \\ & \leq -H \left(M^{(N)}, \text{diag}(\mu_0^{(N)}) A \right) + \frac{n}{2} \log(N), \end{aligned}$$

where the last inequality comes from the second-to-last expression by i) identifying the first term as the KL divergence, ii) noting that the second term is nonpositive, and iii) overestimating the third term by taking $(\mu_0^{(N)})_i = N$ for $i = 1, \dots, n$.

Similarly, underestimating the multinomial coefficients

gives that the log-likelihood can be bounded from below by

$$\begin{aligned}
& \log \left(P_{\mu_0^{(N)}, A} (M^{(N)}) \right) \\
& \geq \sum_{(i,j) \in \mathcal{Z}^{(N)}} \left(m_{ij}^{(N)} \log \left(\frac{(\mu_0^{(N)})_i a_{ij}}{m_{ij}^{(N)}} \right) - \frac{1}{2} \log(m_{ij}^{(N)}) \right) \\
& \quad + \sum_{i \in \mathcal{Y}^{(N)}} \frac{1}{2} \log((\mu_0^{(N)})_i) - \frac{1}{2} n(n-1) \log(2\pi) \\
& \geq -H \left(M^{(N)}, \text{diag}(\mu_0^{(N)}) A \right) \\
& \quad - \frac{1}{2} \left(n^2 + \frac{n(n-1) \log(2\pi)}{\log(N)} \right) \log(N).
\end{aligned}$$

By using the two inequalities, the result follows.

B. Proof of Proposition 2

We solve (10) by a block coordinate ascent in the dual. First note that since μ_0 , A , B , and Φ_t are all elementwise nonnegative, so will the optimal solution $M_{[1:T]}^*$, $D_{[1:T]}^*$, $\mu_{[1:T]}^*$ also be. We can therefore add the constraint $\mu_{[1:T]} \geq 0$ to (10) without changing the optimal solution. For this problem, we relax the constraints in (10) with corresponding dual variables $\lambda_{M_t}, \nu_{M_t}, \lambda_{D_t}, \nu_{D_t}$. Let $\mathfrak{M} := M_{[1:T]}$, $\mathfrak{D} := D_{[1:T]}$, and define the Lagrangian

$$\begin{aligned}
& L(\mathfrak{M}, \mathfrak{D}, \mu_{[1:T]}, \lambda_{\mathfrak{M}}, \nu_{\mathfrak{M}}, \lambda_{\mathfrak{D}}, \nu_{\mathfrak{D}}) \\
& = \sum_{t=1}^T \left(\sum_{ij} m_{ij}^t \log \left(\frac{m_{ij}^t}{\mu_i^{t-1} a_{ij}} \right) + \lambda_{M_t}^T (\mu_{t-1} - M_t \mathbf{1}) \right. \\
& \quad + \nu_{M_t} (\mu_t - M_t^T \mathbf{1}) + \sum_{ij} d_{ij}^t \log \left(\frac{d_{ij}^t}{\mu_i^t b_{ij}} \right) \\
& \quad \left. + \lambda_{D_t}^T (\mu_t - D_t \mathbf{1}) + \nu_{D_t} (\Phi_t - D_t^T \mathbf{1}) \right).
\end{aligned}$$

Minimizing this with respect to the matrices M_t and D_t gives explicit expressions for the optimal solution in terms of $\mu_{[0:T]}$ and the dual variables, i.e.,

$$\begin{aligned}
M_t &= \frac{1}{e} \text{diag}(\mu_{t-1} \odot u_t) A \text{diag}(w_t), \\
D_t &= \frac{1}{e} \text{diag}(\mu_t \odot x_t) B \text{diag}(v_t),
\end{aligned}$$

where $u_t = \exp(\lambda_{M_t})$, $w_t = \exp(\nu_{M_t})$, $x_t = \exp(\lambda_{D_t})$ and $v_t = \exp(\nu_{D_t})$, for $t = 1, \dots, T$. Plugging these into the Lagrangian, we get the modified Lagrangian

$$\begin{aligned}
L(\mu_{[1:T]}, u_{\mathfrak{M}}, w_{\mathfrak{M}}, x_{\mathfrak{D}}, v_{\mathfrak{D}}) &= -\frac{1}{e} \sum_{t=1}^T (\mu_{t-1} \odot u_t)^T A w_t \\
& \quad - \frac{1}{e} \sum_{t=1}^T (\mu_t \odot x_t)^T B v_t + \log(u_1)^T \mu_0 \\
& \quad + \sum_{t=1}^{T-1} \mu_t^T (\log(u_{t+1}) + \log(w_t) + \log(x_t)) \\
& \quad + \mu_T^T (\log(w_T) + \log(x_T)) + \sum_{t=1}^T \log(v_t)^T \Phi_t.
\end{aligned}$$

Noting that since $\mu_{[1:T]}$ occurs linearly in L , for the modified dual functional $\inf_{\mu_{[1:T]} \geq 0} L(\mu_{[1:T]}, \lambda_{\mathfrak{M}}, \nu_{\mathfrak{M}}, \lambda_{\mathfrak{D}}, \nu_{\mathfrak{D}})$ to be bounded from below, the corresponding factors need to be elementwise nonnegative. In this case the corresponding terms will be zero when taking the infimum, and thus the dual problem is to maximize

$$-\frac{1}{e} (\mu_0 \odot u_1)^T A w_1 + \log(u_1)^T \mu_0 + \sum_{t=1}^T \log(v_t)^T \Phi_t \quad (12)$$

subject to

$$\begin{aligned}
& -\frac{1}{e} \text{diag}(u_{t+1}) A w_{t+1} - \frac{1}{e} \text{diag}(x_t) B v_t \\
& \quad + \log(u_{t+1}) + \log(x_t) + \log(w_t) \geq 0
\end{aligned} \quad (13)$$

for $t = 1, \dots, T-1$, and

$$-\frac{1}{e} \text{diag}(x_T) B v_T + \log(x_T) + \log(w_T) \geq 0. \quad (14)$$

Note that neither the objective function (12), nor the first $T-1$ constraints (13) depend on x_T . Thus the optimal choice of x_T is the one creating the most slack in the last constraint, i.e., the one maximizing the first two terms of (14). This is achieved by $x_T = e\mathbf{1} \cdot (Bv_T)$, and for this choice of x_T the constraint (14) can be replaced by

$$-\log(Bv_T) + \log(w_T) \geq 0. \quad (15)$$

Similarly, the slack in (13) is maximized by selecting

$$\begin{aligned}
x_t &= e\mathbf{1} \cdot (Bv_t), \quad t = 1, \dots, T \\
u_t &= e\mathbf{1} \cdot (Aw_t), \quad t = 2, \dots, T.
\end{aligned}$$

and thus the constraints (13) can be replaced by

$$-\log(Aw_{t+1}) - \log(Bv_t) + \log(w_t) \geq 0, \quad (16)$$

for $t = 1, \dots, T-1$. Next, note that if some of the constraints (15) and (16) are not fulfilled with equality, the objective function (12) can be improved by increasing the values of the corresponding v_t . Therefore, in an optimal point we must have that

$$\begin{aligned}
w_T &= Bv_T \\
w_t &= (Bv_t) \odot (Aw_{t+1}) \quad \text{for } t = T-1, \dots, 1.
\end{aligned}$$

This gives an expression for w_1 , that depends on v_t , $t = 1, \dots, T$. Inserting this into the objective (12) leads to an unconstrained problem that depends only on u_1 and v_t , $t = 1, \dots, T$, which we solve using block coordinate ascent. To this end, the unconstrained objective is first maximized with respect to u_1 , which gives

$$u_1 = e\mathbf{1} \cdot (Aw_1).$$

Further, note that the gradient of the unconstrained objective with respect to v_t is

$$-\frac{1}{e} \odot (B^T (y_t \odot Aw_{t+1})) + \Phi_t \cdot / v_t,$$

where y_t is defined by the recursion

$$\begin{aligned}
y_1 &= A^T (\mu_0 \odot u_1), \\
y_t &= A^T (y_{t-1} \odot (Bv_{t-1})), \quad t = 2, \dots, T.
\end{aligned}$$

Hence, maximization with respect to v_t is achieved by

$$v_t = e\Phi_t \cdot (B^T(y_t \odot Aw_{t+1})).$$

As the unconstrained problem is convex and the objective function continuously differentiable, the block coordinate ascent method converges [3, Prop. 2.7.1]. In the limit point, the hidden marginals can be reconstructed from

$$\begin{aligned}\mu_t &= M_t^T \mathbf{1} = \frac{1}{e} \text{diag}(w_t) A^T (\mu_{t-1} \odot u_t) \\ &= \text{diag}(w_t) A^T (\mu_{t-1} \cdot (Aw_t)).\end{aligned}$$

REFERENCES

- [1] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré. Iterative bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.
- [2] G. Bernstein and D. Sheldon. Consistently estimating Markov chains with noisy aggregate data. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pages 1142–1150. Society for Artificial Intelligence and Statistics, 2016.
- [3] D.P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1999.
- [4] S. Blackman and R. Popoli. *Design and analysis of modern tracking systems*. Artech House, Norwood, MA, 1999.
- [5] F. Caravetta and L.B. White. Modelling and estimation for finite state reciprocal processes. *IEEE Transactions on Automatic Control*, 57(9):2190–2202, 2012.
- [6] Y. Chen, G. Conforti, and T. T. Georgiou. Measure-valued spline curves: An optimal transport viewpoint. *SIAM Journal on Mathematical Analysis*, 50(6):5947–5968, 2018.
- [7] Y. Chen, T. Georgiou, and M. Pavon. Entropic and displacement interpolation: A computational approach using the Hilbert metric. *SIAM Journal on Applied Mathematics*, 76(6):2375–2396, 2016.
- [8] Y. Chen, T. Georgiou, M. Pavon, and A. Tannenbaum. Robust transport over networks. *IEEE Transactions on Automatic Control*, 62(9):4675–4682, 2017.
- [9] Y. Chen, T.T. Georgiou, and M. Pavon. On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169(2):671–691, 2016.
- [10] Y. Chen and J. Karlsson. State tracking of linear ensembles via optimal mass transport. *IEEE Control Systems Letters*, 2(2):260–265, 2018.
- [11] T.M. Cover and J.A. Thomas. *Elements of information theory*. John Wiley & Sons, Hoboken, NJ, 2006.
- [12] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2292–2300. Neural Information Processing Systems Foundation, 2013.
- [13] S. Danielsson and J. Flygare. A multi-target graph-constrained HMM localisation approach using sparse wi-fi sensor data. Master thesis, KTH Royal Institute of Technology, 2018.
- [14] D. Dawson, L. Gorostiza, and A. Wakolbinger. Schrödinger processes and large deviations. *Journal of mathematical physics*, 31(10):2385–2388, 1990.
- [15] A. Dembo and O. Zeitouni. *Large deviations techniques and applications*. Springer-Verlag Berlin Heidelberg, 2009.
- [16] R. Ellis. *Entropy, Large Deviations, and Statistical Mechanics*. Springer-Verlag Berlin Heidelberg, 2006.
- [17] F. Elvander, I. Haasler, A. Jakobsson, and J. Karlsson. Tracking and sensor fusion in direction of arrival estimation using optimal mass transport. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 1617–1621. IEEE, 2018.
- [18] M. Fanaswala, V. Krishnamurthy, and L.B. White. Destination-aware target tracking via syntactic signal processing. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3692–3695, 2011.
- [19] H. Föllmer. Random fields and diffusion processes. In P.-L. Hennequin, editor, *École d'Été de Probabilités de Saint-Flour XV–XVII, 1985–87*, volume 1362 of *Lecture Notes in Mathematics*, pages 101–203. Springer, Berlin, Heidelberg, 1988.
- [20] T.T. Georgiou and M. Pavon. Positive contraction mappings for classical and quantum Schrödinger systems. *Journal of Mathematical Physics*, 56(3):033301, 2015.
- [21] Z. Ghahramani. An introduction to hidden Markov models and Bayesian networks. *International journal of pattern recognition and artificial intelligence*, 15(1):9–42, 2001.
- [22] B. Jamison. Reciprocal processes. *Probability Theory and Related Fields*, 30(1):65–86, 1974.
- [23] J. Karlsson and A. Ringh. Generalized Sinkhorn iterations for regularizing inverse problems using optimal mass transport. *SIAM Journal on Imaging Sciences* 10(4), 1935–1962, 2017.
- [24] G. King. *A Solution to the Ecological Inference Problem: Reconstructing Individual Behavior from Aggregate Data*. Princeton University Press, Princeton, NJ, 1997.
- [25] B.C. Levy, R. Frezza, and A.J. Krener. Modeling and estimation of discrete-time gaussian reciprocal processes. *IEEE Transactions on Automatic Control*, 35(9):1013–1023, 1990.
- [26] C. Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete & Continuous Dynamical Systems - A*, 34(4):1533–1574, 2014.
- [27] T. Mikami. Monge's problem with a quadratic cost by the zero-noise limit of h -path processes. *Probability Theory and Related Fields*, 129(2):245–260, 2004.
- [28] M. Pavon and F. Ticozzi. Discrete-time classical and quantum Markovian evolutions: Maximum entropy problems on path space. *Journal of Mathematical Physics*, 51(4):042104, 2010.
- [29] L.R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [30] H. Robbins. A remark on Stirling's formula. *The American Mathematical Monthly*, 62(1):26–29, 1955.
- [31] E. Schrödinger. Über die Umkehrung der Naturgesetze. *Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-mathematische Klasse*, pages 144–153, 1931.
- [32] G. Stamatescu, L.B. White, and R. Bruce-Doust. Track extraction with hidden reciprocal chains. *IEEE Transactions on Automatic Control*, 63(4):1097–1104, 2018.
- [33] C. Villani. *Optimal transport: Old and new*. Springer, Berlin Heidelberg, 2008.
- [34] L.B. White and H.X. Vu. Maximum likelihood sequence estimation for hidden reciprocal processes. *IEEE Transactions on Automatic Control*, 58(10):2670–2674, 2013.
- [35] S. Zeng, S. Waldherr, C. Ebenbauer, and F. Allgöwer. Ensemble observability of linear systems. *IEEE Transactions on Automatic Control*, 61(6):1452–1465, 2016.