Hidden Markov modeling of single particle diffusion with stochastic tethering

Amit Federbush,[1,2] Amit Moscovich,[3] and Yohai Bar-Sinai[1,2]

[1]*Department of Condensed Matter Physics, Tel Aviv University, Ramat Aviv, Tel Aviv, 69978, Israel*
[2]*The Center for Physics and Chemistry of Living Systems, Tel Aviv University, Tel Aviv 69978, Israel*
[3]*Department of Statistics and Operations Research, Tel Aviv University, 69978, Israel*
(Dated: August 3, 2023)

The statistics of the diffusive motion of particles often serve as an experimental proxy for their interaction with the environment. However, inferring the physical properties from the observed trajectories is challenging. Inspired by a recent experiment, here we analyze the problem of particles undergoing two-dimensional Brownian motion with transient tethering to the surface. We model the problem as a Hidden Markov Model where the physical position is observed, and the tethering state is hidden. We develop an alternating maximization algorithm to infer the hidden state of the particle and estimate the physical parameters of the system. The crux of our method is a saddle-point-like approximation, which involves finding the most likely sequence of hidden states and estimating the physical parameters from it. Extensive numerical tests demonstrate that our algorithm reliably finds the model parameters, and is insensitive to the initial guess. We discuss the different regimes of physical parameters and the algorithm's performance in these regimes. We also provide a ready-to-use open source implementation of our algorithm.

I. INTRODUCTION

Since the early days of statistical mechanics, the statistics of the stochastic motion of mesoscopic particles are an important experimental probe for their microscopic properties. Most prominently, the Gaussian statistics of Brownian motion provided an experimental proof of the atomic nature of matter [1–3] and was used to measure Avogadro's number [4]. To this day, new models are actively developed to explain deviations from purely Brownian statistics in biological and colloidal systems [5, 6]. Of specific interest are systems undergoing anomalous diffusion which exhibit a Fickian behavior, i.e. a mean-squared-displacement linear with time, yet a non-Gaussian behavior of the displacement statistics [5–7]. These deviations from Gaussianity have been suggested as an accessible probe for various experiments [8–13].

In the analysis of such systems, two key sources for non-Gaussianity are considered. First, particles may undergo several different types of diffusion modes, stochastically switching between them [14–19]. Second, particles may be transiently confined to a small region [9, 20–29]. In this work, inspired by the experimental system of Chakraborty et. al. [9], we focus on the latter, and specifically on two-dimensional (2D) diffusion with transient tethering to the underlying surface. Fig. 1 depicts an example trajectory from such an experiment. In such experiments, colloids or nano-particles are coated with molecules of interest (typically, peptides) and undergo 2D diffusion on a surface coated with a different molecule. The interaction between the molecules leads to stochastic transient tethering of the particles to the surface. The experiment aims to extract information regarding the interaction between the peptides from the frequency of these tethering and untethering events.

However, identifying these events may be challenging, since the tethering to the surface is not directly observed, but rather needs to be inferred from the slowdown in the observed trajectories. Several computational methods have been suggested to tackle this in the past [22, 24, 29]. In this work, we present a simple and computationally efficient algorithm to infer tethering ad untethering events from observed trajectories, to estimate the tethering and untethering rates, the diffusion coefficient, and the effective confinement area of the interaction potential.

The structure of this manuscript is as follows: In Section II we describe the stochastic Hidden Markov Model which we use to model the problem, discuss the different time scales that the model introduces, and define the regimes of model parameters where our methods are valid. In Section III we describe our algorithm and its underlying approximations, and in Section IV we present our numerical results on synthetic data. Finally, in Section V we discuss the results and possible outlooks. The full source code of the implemented algorithm is available on GitHub [30].

II. STOCHASTIC MODEL

Our Markov model describes a particle alternating between a freely diffusing state and a tethered state in 2D. The transition between the states is modeled as a standard two-state continuous-time Markov chain with characteristic times τ_0 and τ_1, which are the inverses of the average rates of tethering and untethering events, respectively. We denote the state of the particle by $S(t)$ where $S = 0$ corresponds to the free state and $S = 1$ to the tethered state. In the free state the particle undergoes standard Brownian motion, while in the tethered state it is also confined to a harmonic potential centered at the tether point $X^*(t)$. Explicitly, the position of the particle $X(t)$ follows an Ornstein-Uhlenbeck process with a
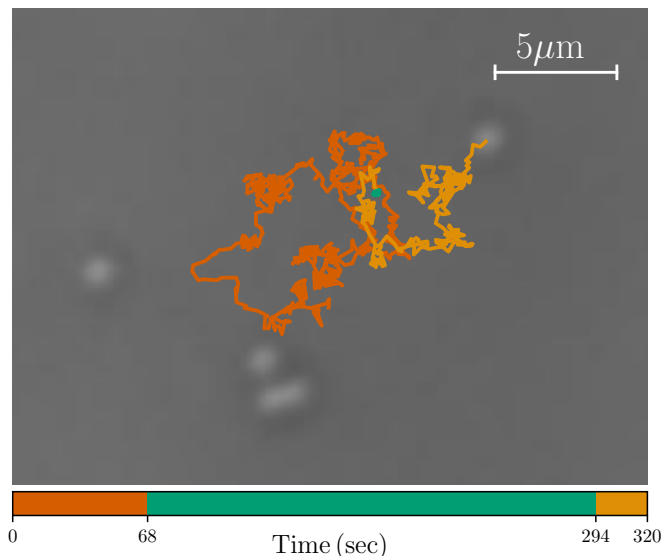
FIG. 1. A snapshot of an experimental microscopy video of peptide-coated microparticles, in a system similar to that of Chakraborty et. al. [9]. The colored curve represents the top-right particle's recorded trajectory in the 320 seconds preceding the snapshot. The trajectory is colored based on time, as indicated by the timeline at the bottom. During the time segment $t \in [68, 294]$ the particle is confined to a small region of space (shown in green) due to tethering. Image courtesy of Roy Beck and Amandeep Sekhon.

state-dependent potential:

$$\dot{X}(t) = -\frac{k}{\gamma}(X(t) - X^*(t))S(t) + \sqrt{2D}\xi(t) \qquad (1)$$

where k is the spring constant, γ is the friction coefficient, $S(t) \in \{0,1\}$ indicates the free/tethered state, D is the diffusion coefficient, and ξ is a standard two-dimensional Wiener process $\langle \xi_i(t)\xi_j(t') \rangle = \delta_{ij}\delta(t - t')$. Observe that only the ratio k/γ plays a role in the model, rather than k and γ individually. We assume the particle is tethered to the point at which the transition $S = 0 \to 1$ occurred. That is, if the particle became tethered at time t_1, and stayed tethered until it untethered at $t_2 > t_1$, then

$$X^*(t) = X(t_1), \qquad t \in [t_1, \ t_2). \qquad (2)$$

To summarize, the particle dynamics are modeled by a continuous-time Markov process $F(t) = (X(t), S(t), X^*(t))$. The model is specified by 4 parameters: τ_0, τ_1, D and the ratio k/γ. It is more convenient to work with the variable $A = (D\gamma)/k$ which is the characteristic area that the particle explores in the tethered state. For notational purposes, we group the model parameters as $\Theta := (\tau_0, \tau_1, D, A)$. Lastly, we note that $\frac{A}{D} = \frac{\gamma}{k}$ has units of time. Its meaning will be discussed below.

The transition probabilities of S are Poissonian and are well known. In the untethered state $S(t) = 0$, the particle undergoes classic Brownian motion. Hence its

position at $t + \Delta t$ is normally distributed around $X(t)$,

$$P\left(X(t + \Delta t)\right) = \frac{1}{4\pi D \Delta t} \exp\left[\frac{-\left(X(t + \Delta t) - X(t)\right)^2}{4D\Delta t}\right] \tag{3}$$

In the tethered state $S(t) = 1$, the probability density function of $X(t + \Delta t)$ is given by the solution of the Fokker-Planck equation with strong friction in a harmonic potential around the anchor point $X^*(t)$ [31].

$$P\left(X(t + \Delta t)\right) = \tag{4}$$

$$\frac{1}{2\pi A'(\Delta t)} \exp\left[-\frac{\left(X(t + \Delta t) - \phi(\Delta t)X(t) - (1 - \phi(\Delta t))X^*(t)\right)^2}{2A'(\Delta t)}\right]$$

where we defined two auxiliary variables:

$$\phi(\Delta t) = \exp\left(-\frac{D\Delta t}{A}\right), \quad A'(\Delta t) = (1 - \phi^2(\Delta t))A. \tag{5}$$

A. Discretized dynamics

The model defined above describes the full continuous-time dynamics. In principle, we could use it to estimate the model parameters $\Theta = (\tau_0, \tau_1, D, A)$. However, the experimental setup poses two difficulties: First, we can only measure the particle positions $X(t)$ and do not have access to the particle states $S(t)$ and tether points $X^*(t)$. This is the core challenge of the problem which we address in Section III. Second, we only sample the process at discrete times $t_1, t_2, \ldots, t_N$, separated by a finite time resolution $\Delta t = t_i - t_{i-1}$. To account for this, we define a discrete Markov process analogue of the continuous process, where the tether point is constrained to be one of the previously observed positions. This discrete Markov process defines a sequence of random states $\{F_n\}_{n=1}^N$ where

$$F_n := F(t_n) = (X(t_n), S(t_n), X^*(t_n)). \tag{6}$$

The transition probabilities are given by the product of the transition probabilities of S and X:

$$P(F_{n+1}|F_n) = P(S_{n+1}|S_n; \Theta)P(X_{n+1}|F_n; \Theta) \tag{7}$$

The state transitions are Poissonian and to leading order in Δt read

$$P\left(S_{n+1}|S_n\right) = \begin{cases} 1 - \frac{\Delta t}{\tau_n} & S_{n+1} = S_n \\ \frac{\Delta t}{\tau_n} & S_{n+1} \neq S_n, \end{cases} \tag{8}$$

where $\tau_n = \tau_0$ if $S_n = 0$ and $\tau_n = \tau_1$ if $S_n = 1$. The distribution of the particle position at the next step follows Eq. (3) and (4),

$$P\left(X_{n+1}|F_n; \Theta\right) = \tag{9}$$

$$\begin{cases} \frac{1}{4\pi D \Delta t} e^{-\frac{1}{4D\Delta t}\left(X_{n+1} - X_n\right)^2} & S_n = 0 \\ \frac{1}{2\pi A'} e^{-\frac{1}{2A'}\left(X_{n+1} - \phi X_n - (1-\phi)X_n^*\right)^2} & S_n = 1 \end{cases}$$
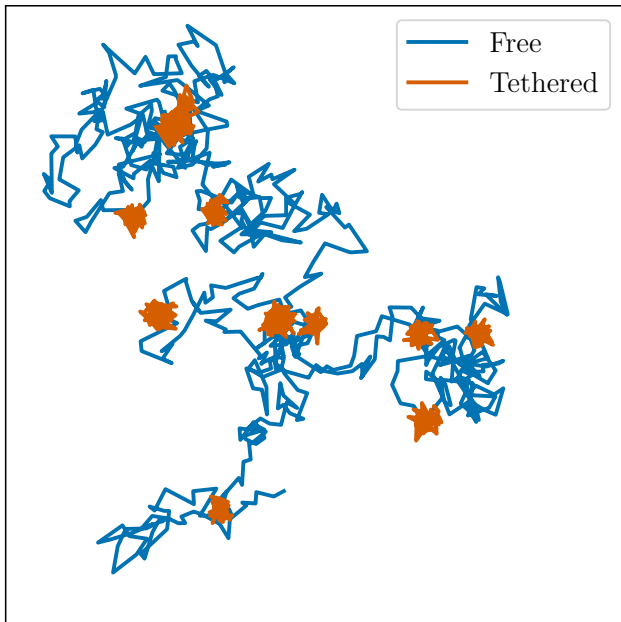
FIG. 2. An example trajectory of $N = 2000$ steps generated from the discrete model with $\tau_0 = \tau_1 = 100$, $D = 1$ and $A = 0.5$. The time step is $\Delta t = 1$. The trajectory is color-coded according to the particle's state, where blue and orange depict the free and tethered states, respectively.

where now ϕ and A' are constants that depend on A, D and the time step Δt as in Eq. (5). The discretization of Eq. (2) reads

$$X_{n+1}^* = \begin{cases} X_{n+1} & S_{n+1} = 1 \text{ and } S_n = 0 \\ X_n^* & \text{otherwise} \end{cases} \quad (10)$$

Furthermore, we assume that at the beginning of the measurement the particle is either free ($S_1 = 0$) or tethered at X_1 ($S_1 = 1, X_1^* = X_1$) . Importantly, X_n^* can only attain values among the (finitely many) previously visited positions $X_1, \ldots, X_n$, and the sequence of tether points $\{X_n^*\}$ is completely determined by the history of $\{X_n\}$ and $\{S_n\}$.

As usual for Markov processes, the probability of a trajectory $\{F_n\}_{n=1}^N$ is the product of the transition probabilities between each of its states, thus the log-probability is additive:

$$\log P\left(\{F_n\}_{n=1}^N; \Theta\right)$$
$$= \log P(F_1) + \sum_{n=1}^{N-1} \log P(F_{n+1}|F_n; \Theta). \quad (11)$$

An example trajectory generated from the discrete model is depicted in Fig. 2.

B. Time scales

As seen above, the problem involves multiple time scales. Three of them are given by the model parameters and represent the underlying physics: τ_0, τ_1, and $\frac{A}{D}$ which is the equilibration time of the Brownian particle with the harmonic potential. The other two time scales, Δt and the total sampling time $T = t_N - t_1$ are properties of the experiment.

We employ several realistic working assumptions about these time scales that greatly reduce computational complexity. First, we assume that $\tau_0, \tau_1 \gg \frac{A}{D}$, which physically means that the rate of tethering/untethering events is much smaller than the inverse of the harmonic equilibration time of the particle. Violation of this condition means that the particle can untether before the tethering potential has a significant effect and tethering events will not be experimentally discernible. This regime corresponds to standard Brownian motion with an effective diffusion constant smaller than D.

Second, we require $\Delta t \ll \tau_0, \tau_1$, to avoid the possibility of multiple transitions of tethering or untethering events within a single sampling interval. This assumption is experimentally realistic since modern cameras can easily achieve frame rates larger than 10^3Hz, and in many experiments, τ_0 and τ_1 are of orders of at least seconds [9, 20, 25]. In any case, if $\Delta t \approx \tau_0, \tau_1$, this would again correspond to an essentially pure diffusive behavior in the discretized data.

Third, if $\Delta t < \frac{A}{D}$, the time discretization resolves the equilibration of a tethered particle with its confining potential. Since our goal is only to extract the physical parameters of the system, such resolution does not add relevant information and only increases the computational cost. Therefore, under-sampling the discrete dynamics to increase Δt should not lead to a significant loss of accuracy in the estimation of τ_0, τ_1 but would greatly reduce the search space. This will be explicitly demonstrated in Section IV. Even if the trajectories are experimentally measured with small Δt, we can safely undersample them such that $\frac{A}{D} \lesssim \Delta t$, or in other words $\phi \lesssim e^{-1}$, cf. Eq. (5).

For concreteness, we mention the experimental parameters of the diffusing nanoparticles system of Chakraborty et. al. [9]. In this system, $\tau_0, \tau_1 \sim 1$s, $D \sim 10\mu\text{m}^2/\text{s}$, $A \sim 1\mu\text{m}^2$, which means a sampling interval of $\Delta t = A/D = 0.1$s is sufficient and is easily achievable experimentally.

To conclude this discussion, we assume the following separation of time scales:

$$\frac{A}{D} \lesssim \Delta t \ll \tau_0, \tau_1, \quad (12)$$

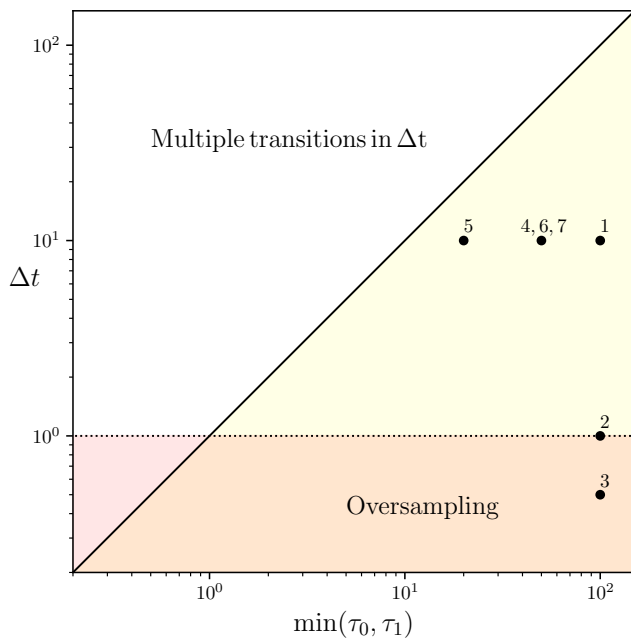Using these assumptions, the leading-order expansion of Eq. (9), which is first order in Δt and zeroth order in ϕ,

FIG. 3. A section of the phase space of $\Delta t, \tau_0, \tau_1$, in time units of $\frac{A}{D}$. Note the logarithmic scale. The yellow region is where our assumptions hold. The bottom horizontal boundary of the yellow region corresponds to the inequality $\frac{A}{D} \lesssim \Delta t$ and the diagonal boundary corresponds to the inequality $\Delta t \ll \tau_0, \tau_1$. The bottom and left regions are annotated according to the discussion in this subsection. The markers correspond to the seven regimes to be analyzed in Section IV.

reads

$$P\left(S_{n+1}|S_n\right) = \begin{cases} 1 - \frac{\Delta t}{\tau_n} & S_{n+1} = S_n \\ \frac{\Delta t}{\tau_n} & S_{n+1} \neq S_n \end{cases}, \qquad (13)$$

$$P\left(X_{n+1}|F_n\right) = \begin{cases} \frac{1}{4\pi D \Delta t} e^{-\frac{1}{4D\Delta t}\left(X_{n+1} - X_n\right)^2} & S_n = 0 \\ \frac{1}{2\pi A} e^{-\frac{1}{2A}\left(X_{n+1} - X_n^*\right)^2} & S_n = 1 \end{cases}$$

where $\tau_n = \tau_0$ if $S_n = 0$ and $\tau_n = \tau_1$ if $S_n = 1$.

To summarize the discussion regarding time scales, Fig. 3 illustrates the section of phase space where our assumptions hold (in time units of $\frac{A}{D}$). For a point in the yellow region in the figure, the farther it is from the two boundary lines, the better is the separation of scales and the better our assumptions hold. The marked points correspond to regimes that will be analyzed in Section IV.

III.   OUR METHOD

Our problem is as follows: in an experiment, we can measure the observed states $\{X_n\}$, but we do not have access to the hidden states $\{S_n, X_n^*\}$, nor to the model parameters $\Theta = (\tau_0, \tau_1, D, A)$. We use the term *hidden path* to denote the sequence of the hidden states $\{S_n\}_{n=1}^N$

in time from which the sequence $\{X_n^*\}_{n=1}^N$ can be determined. The goal is to infer the model parameters Θ from a series of measurements. To this end, we developed an alternating maximization algorithm [32], similar to the classical expectation-maximization (EM) algorithm [33], to estimate both the hidden states and the parameters of our hidden Markov model.

If the hidden path is known, the problem of optimal parameter estimation is fairly standard and amounts to maximizing the likelihood of the model parameters, which according to Bayes' rule is proportional to the exponential of Eq. (11). Maximizing $P(\Theta|\{F_n\})$ can be done numerically or using analytical approximations, as described below. However, when the hidden path is not known, the likelihood function to consider is

$$\mathcal{L}\left(\Theta|\{X_n\}\right) = \frac{P(\Theta)}{P(\{X_n\})} P\left(\{X_n\}|\Theta\right)$$
$$= \frac{P(\Theta)}{P(\{X_n\})} \sum_{\{S_n\}} P\left(\{F_n\}|\Theta\right), \qquad (14)$$

where the sum is over all 2^N possibilities for the hidden paths, $P(\Theta)$ is the prior probability distribution for the model parameters, and the evidence $P(\{X_n\})$ is a constant we may ignore [34]. Recall that $\{F_n\} = \{X_n, S_n, X_n^*\}$ is the trajectory of both the observed and hidden states of the particle, and the log-probability of such a trajectory is the sum of the log-probability of each step, as in Eq. (11). Computing the sum in Eq. (14) is intractable for typical values of N. However, numerical evidence shows that most hidden paths are very unlikely and thus have a negligible contribution to this sum. For more on this, see Section IV C. Taking a uniform prior, $P(\Theta) = \mathrm{const}$, we posit that:

$$\log \mathcal{L}(\Theta|\{X_n\}) \sim \max_{\{S_n\}} \log P(\{F_n\}|\Theta). \qquad (15)$$

A key idea of our method is to use the RHS of Eq. (15) as a proxy for the computationally intractable LHS. It is analogous to the saddle-point approximation from statistical mechanics, where the integral is replaced with the maximum of the integrand. To compute the RHS of Eq. (15), we need only to find the most likely hidden path $\{\hat{S}_n\}_{n=1}^N$, given the model parameters Θ. This discrete optimization problem can be efficiently solved using the Viterbi algorithm from dynamic programming[34, 35]. Below we briefly describe the parameter estimation method and how the Viterbi algorithm can be implemented for our model. Then we present our alternating maximization approach for estimating the maximum likelihood model parameters.

A.   Parameter estimation

Given the most likely hidden path $\{\hat{S}_n\}$, maximizing the log-likelihood in Eq. (15) is maximizing a sum of the

log terms from Eq. (13). We use the following maximum-likelihood estimators (MLE) to approximate the most likely model parameters:

$$\hat{\tau}_0 := \frac{N_{00} + N_{01}}{N_{01}} \Delta t \qquad \hat{\tau}_1 := \frac{N_{11} + N_{10}}{N_{10}} \Delta t$$

$$\hat{D} := \frac{1}{4(N_{00} + N_{01})\Delta t} \sum_{n=1}^{N-1} (1 - \hat{S}_n)(X_{n+1} - X_n)^2 \quad (16)$$

$$\hat{A} := \frac{1}{2(N_{10} + N_{11})} \sum_{n=1}^{N-1} \hat{S}_n (X_{n+1} - X_n^*)^2,$$

where N_{ij} is the number of $i \to j$ transitions of $\hat{S}_n$ that occur along the trajectory. The derivation of these estimators is trivial by taking the derivative of Eq. (15) with respect to each model parameter and equating to zero. We group the four MLEs as $\hat{\Theta} := (\hat{\tau}_0, \hat{\tau}_1, \hat{D}, \hat{A})$.

B. Finding the most likely sequence of states

We now describe how to find the maximum likelihood path $\{\hat{S}_n\}$, conditioned on the model parameters Θ. Recall that the tether point X_n^* can only assume one of the previously visited positions $X_1, X_2, ..., X_n$, cf. Eq. (10). We represent the set of all possible paths as a directed layer graph (Fig. 4), known as a *trellis* in the Viterbi literature. In this graph, columns correspond to the discrete time n and rows to the tethered state X_n^* which can be either one of $X_1, ..., X_n$ or *free* (untethered). Each vertex represents the state of the particle at time n and each full path from left to right corresponds to a specific sequence of tethered/free states.

We set the edge weights to be the log of the transition probabilities between states according to Eq. (7). With this choice of edge weights, the log-likelihood of a path (Eq. (11)) is just $\log P(F_1)$ plus the sum of edge weights along the corresponding path in the graph. Thus, finding the maximum-likelihood sequence of hidden states is reduced to the problem of finding the maximizing path in a directed layer graph. The latter problem is efficiently solved using the Viterbi algorithm [35]. This algorithm scans the trellis column by column from left to right and computes, for each vertex, the maximum-weight path that ends at that vertex.

C. The alternating maximization algorithm

Given an estimate of the hidden path $\{S_n\}$, we can apply Eq. (16) to obtain the maximum likelihood estimate of the model parameters $\Theta = (\tau_0, \tau_1, D, A)$. Conversely, given an estimate of the model parameters Θ, it is easy to find the maximum likelihood hidden path using the Viterbi algorithm as explained in the previous subsection. Combining these two observations naturally leads to a fast alternating maximization procedure for
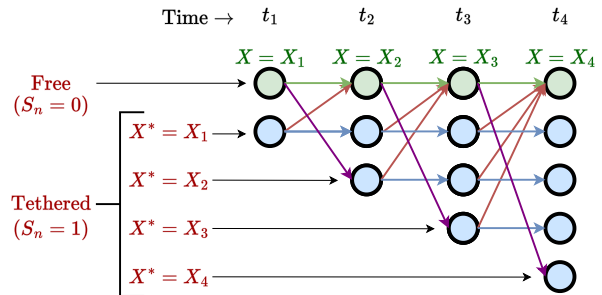

FIG. 4. A trellis graph of the observed and hidden states, given the observed particle positions $X_1, ..., X_4$. Each node represents a different state $F_n = (X_n, S_n, X_n^*)$, with the row corresponding to the hidden state, and the column to the time n and the observed state X_n. The edges represent allowed state transitions and are weighted according to the logarithm of Eq. (7). Each trajectory $\{F_n\}$ is given as a path on the graph that advances from left to right.

estimating both the most likely model parameters $\hat{\Theta}$ and the most likely hidden path:

1. **Initial guess**: Guess an initial value $\Theta^{(0)}$ for the four model parameters.

2. **Path maximization step**: Conditioned on the current parameter estimate $\Theta^{(m)}$, apply the Viterbi algorithm to find $\{S_n^{(m)}\}_{n=1}^N$, the most likely hidden path given $\Theta^{(m)}$,
$$S^{(m)} := \arg\max_{\{S_n\}} \mathcal{L}(\{S_n\} | \Theta^{(m)}, \{X_n\}). \qquad (17)$$

3. **Parameter maximization step**: Use Eq. (16) to obtain the maximum likelihood estimate of the model parameters conditioned on the current estimate of the hidden path,
$$\Theta^{(m+1)} := \arg\max_{\Theta} \mathcal{L}(\Theta | S^{(m)}, \{X_n\}). \qquad (18)$$

4. **Convergence**: Repeat steps 2-3 until $\frac{\left| \theta^{(m+1)} - \theta^{(m)} \right|}{\theta^{(m)}} \leq \epsilon$ for each $\theta \in \Theta$.

Since the sample space is discrete, convergence typically occurs exactly, i.e. $\Theta^{(n+1)} = \Theta^{(n)}$. However, in our experiments we used $\epsilon = 10^{-3}$ to stop the iterations when the relative change is small. The specific threshold 10^{-3} is inconsequential. Furthermore, we set the maximum number of iterations to 20, to prevent the possibility of infinite loops that alternate between several discrete hidden paths without satisfying the convergence criterion (in practice, less than 10 iterations typically suffice for convergence). We also define a criterion for divergence: if after the parameter maximization step the estimators $\hat{\tau}_0$ or $\hat{\tau}_1$ exceed the arbitrary threshold of $0.9T$, we say the algorithm diverged and stop the iterations. This is done for simplicity since we wish to avoid the cases where the particle stays in one tether state for the entire sampling

time, because then some of the MLEs diverge. In practice, we test the regimes where $T \gg \tau_0, \tau_1$ so most runs do not diverge. To summarize, we say the algorithm has converged if the convergence criterion was satisfied without first triggering the divergence criterion. In addition to the above, our implementation of the Viterbi algorithm includes pruning some of the trellis edges (cf. Fig. 4) to improve the running time. The Viterbi algorithm scans each edge on the graph once, leading to an $O(N^2)$ complexity. However most paths are not likely, e.g. a path where at some point $S_n = 1$ and $|X_n - X_n^*|^2 \gg A$, meaning the particle roams far from its tether point. We implement the pruning by discarding the paths of all but the q most likely tethered nodes at each time step (column) during the algorithm's execution. This reduces the complexity to $O(qN)$. We used $q = 10$ in our tests, and we numerically tested this value and found no significant impact on the results compared to no pruning (where $q = N$).

IV.   RESULTS

We test the algorithm's performance on synthetically generated trajectories, whose model parameters Θ are known. The implementation of the algorithm in Python, as well as the code to generate all figures in this manuscript, is available at [30]. In all trajectories we use $D = 1$ and $A = 1$ and a total sampling time of $T = 10000$. Note that both D and A can always be set to unity by properly choosing length and time units, so this choice is without loss of generality. Thus, Θ is specified by three scalars: the sampling time Δt and the model parameters τ_0, τ_1, all expressed in time units of $\frac{A}{D} = 1$. We focus on 7 different regimes of Θ, detailed in Table I. We mention that the conditions in Eq. (12) are barely satisfied in regimes 3 and 5. This is intentional as we wish to test the algorithm slightly beyond its bounds.

A.   Stability

As described above, the algorithm requires an initial guess of the model parameters $\Theta^{(0)}$. We found empirically that the algorithm converges to a single fixed point almost regardless of the initial parameters. To demonstrate this, we show in Fig. 5 the results of the algorithm when applied to a single observed trajectory, with 1000 different $\Theta^{(0)}$'s, randomly drawn from a log-uniform distribution spanning two decades around Θ. It is seen that the algorithm strongly converges towards a single fixed point of the parameter estimators. We mention this test's wide prior distribution for $\Theta^{(0)}$ is intended for illustrating the algorithm's stability, and in a practical situation, better priors can often be used, especially for D and A.

The convergence to a fixed point is similar in all seven parameter regimes we tested, however the figure displays the convergence for a trajectory corresponding to regime

1 (cf. Table I). Out of the 1000 algorithm runs, 960 converged and only these runs were taken into account in the figure. The other runs diverged because some initial parameters led to a trajectory being labeled as entirely free. Since the outcome of the algorithm is largely independent of the initial guess, in what follows we use the true model parameters Θ as the initial guess $\Theta^{(0)}$.

B.   Recovering the parameters and the hidden states

For each of the regimes in Table I, we generated 1000 synthetic trajectories and ran the algorithm with initial parameters that are equal to the true model parameters, as discussed in the previous subsection regarding stability. Over 98% of the runs in each regime have converged (the other runs entered a loop and stopped after the maximum number of iterations). Runs that converge typically do so within 3-8 iterations.

We define the accuracy as the fraction of time steps for which the algorithm predicted both the correct state S_n and tether point X_n^* (if $S_n = 1$). The mean and standard deviation of the accuracy over all converged runs are detailed in Table I. It is seen that the accuracy is fairly high in all regimes, and that it decreases as Δt approaches $\frac{A}{D}$ from above or τ_0, τ_1 from below, consistently with the constraints of Eq. (12).

Next, we examine how well the algorithm recovers the model parameters. For each regime, Table I depicts the mean and the standard deviation of the algorithm's estimates $\hat{\tau}_0, \hat{\tau}_1, \hat{D}, \hat{A}$ (over all converged runs). These results exhibit two clear trends: the temporal parameters τ_0, τ_1 are consistently overestimated, while the spatial ones D, A are correctly estimated (although there is a very slight yet consistent underestimate of A in most regimes). The overestimate of the temporal parameters ranges from 20% to 120%.

To investigate the effect of the sampling time Δt, we focus on the three regimes $1, 2, 3$ which differ only by Δt, as detailed in Table I. The histograms of the four model parameters are depicted in Fig. 6 (as KDE plots). It is seen that there is no significant effect for reducing Δt from 10 to 0.5, with the parameter estimates being similar (despite the accuracy decreasing as Δt decreases). These results support the claim described in Section II that the sampling time can be increased without losing information, as long as Eq. (12) holds. We mention that the width of the distribution for the spatial parameters D, A decreases with Δt. This is expected as the number of samples N increases as Δt decreases when T is fixed.

By focusing on regimes 1, 4, and 5, the effect of τ_0, τ_1, when they are equal, can be isolated. As τ_0 and τ_1 decrease, the accuracy decreases and the relative overestimate of τ_0, τ_1 increases. This is consistent with Eq. (12), since decreasing τ_0, τ_1 challenges the assumption that $\Delta t \ll \tau_0, \tau_1$.

Finally, regimes 6 and 7 depict asymmetry in τ_0, τ_1. It

Regime	D, A	Δt	τ_0	τ_1	Acc. (%)	$\hat{\tau}_0$	$\hat{\tau}_1$	$\hat{D}$	$\hat{A}$
1	1	10	100	100	96 ± 2	131 ± 24	130 ± 19	1.00 ± 0.05	0.99 ± 0.05
2	1	1	100	100	94 ± 2	122 ± 21	122 ± 16	1.00 ± 0.01	0.99 ± 0.02
3	1	0.5	100	100	88 ± 4	125 ± 22	123 ± 17	1.00 ± 0.01	0.99 ± 0.02
4	1	10	50	50	93 ± 2	77 ± 11	75 ± 8	0.99 ± 0.05	0.98 ± 0.05
5	1	10	20	20	87 ± 2	47 ± 9	43 ± 5	0.97 ± 0.06	0.94 ± 0.06
6	1	10	200	50	96 ± 1	356 ± 95	79 ± 12	0.99 ± 0.04	0.97 ± 0.09
7	1	10	50	200	97 ± 2	60 ± 11	248 ± 43	1.01 ± 0.08	1.00 ± 0.04

TABLE I. Model parameters, accuracy, and estimated parameters for the 7 regimes analyzed. The accuracy is defined as the fraction of time frames at which the algorithm correctly predicted the tethered state S and the tether point X^*.
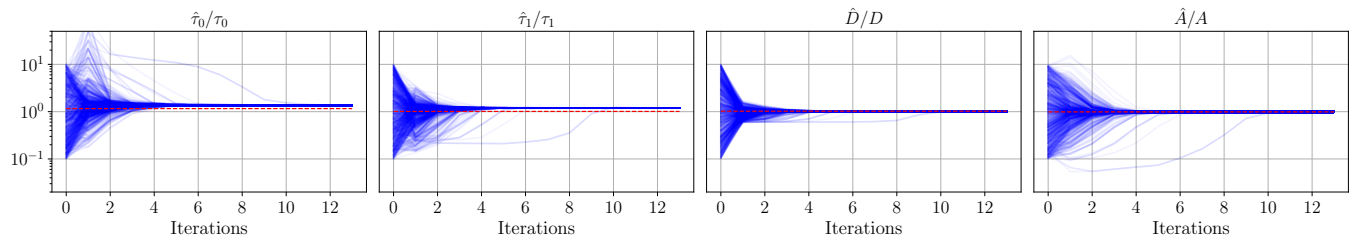

FIG. 5. **The algorithm is insensitive to the initialization of $\Theta^{(0)}$:** Here we show the evolution of the estimators $\hat{\tau}_0, \hat{\tau}_1, \hat{D}, \hat{A}$ of the four model parameters along the algorithm iterations over a single synthetic trajectory. The values are normalized by the true parameter values so that a perfect prediction would be unity in all panels. The first iteration is the initial parameter guess. We drew 1000 initial guesses at random from a log-uniform distribution around the true value and ran the algorithm, keeping only the 960 runs that converged. The evolution of the estimators of each (converged) run is shown, exhibiting a strong convergence towards a single fixed point. Note that the vertical axis is logarithmic. The dashed red line represents the most likely parameter value, given the true hidden path. It does not coincide, in general, with the true model parameters.

is seen that when $\tau_0 < \tau_1$, i.e. when the particle typically spends more time tethered than free, the parameter estimates are better, with less overestimation for the temporal parameters, compared to the opposite case.

C. K most likely paths

Our saddle-point approximation in Eq. (15) replaces the sum over all hidden paths with the likelihood value of the single most likely path. An immediate generalization is to sum over the top K likeliest paths, which would improve the estimation accuracy of $\mathcal{L}$. Given the top K likeliest paths, we can estimate the model parameters as a weighted sum of the MLEs of each path, cf. Eq. (16) as

$$\hat{\Theta}^{[K]} = \frac{\sum_{i=1}^{K} P(\{F_n\}^{[i]}|\Theta)\hat{\Theta}(\{F_n\}^{[i]})}{\sum_{i=1}^{K} P(\{F_n\}^{[i]}|\Theta)}, \qquad (19)$$

where $\{F_n\}^{[i]}$ is the trajectory corresponding to the ith most likely hidden path. Finding the K highest likelihood paths can be easily done by a slight modification of the Viterbi algorithm, also known as list Viterbi [36–38]. The modification is discarding all but the K most likely paths ending at each node, rather than all but the most likely path. Doing this (with our naive implementation)

increases the computational complexity by K^2.

In our numerical tests we found no significant improvement in the results when increasing K up to $K = 100$. This is because the K most likely paths are very similar, differing from each other in just a few steps, leading to very similar MLEs of Θ. Typically, the top paths only differ by slight perturbations of the tethering and untethering times and do not display qualitative differences.

To demonstrate this, Fig. 7 depicts the true hidden path of a single trajectory, and the top $K = 10$ most likely hidden paths given the "oracle parameters" which are the most likely parameters given the true hidden path (similar to the red dashed lines in Fig. 5), and the top $K = 10$ most likely hidden paths given the estimated parameters, i.e. the algorithm's output. It is seen that the most likely hidden paths are all qualitatively very similar. They are also similar to the true hidden path, except for the brief tethered interval at time $t = 93$, and the brief free interval at time $t = 220$, which appear in the true hidden path but not in the most likely hidden paths.

Therefore, because the results were similar for K up to 100, in the previous section we presented only results pertaining to $K = 1$. However, summing over the top K paths might be useful in different parameter regimes.
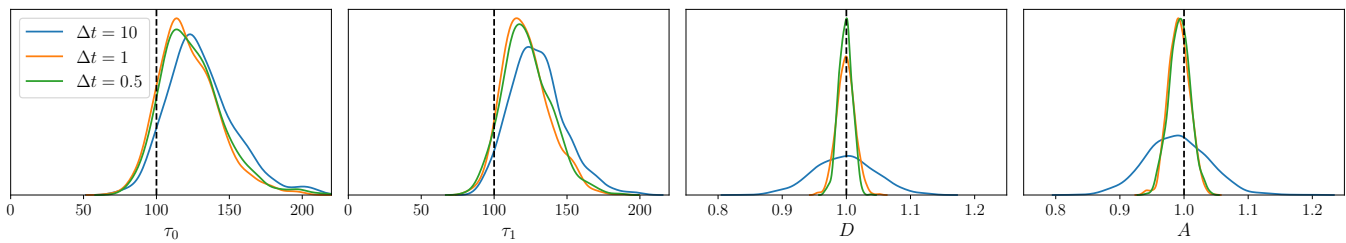
FIG. 6. KDE plots of the algorithm's output parameter values for regimes 1,2,3 with $\tau_0 = \tau_1 = 100$ and $\Delta t = 10, 1, 0.5$, respectively. The dashed lines represent the true parameter values.
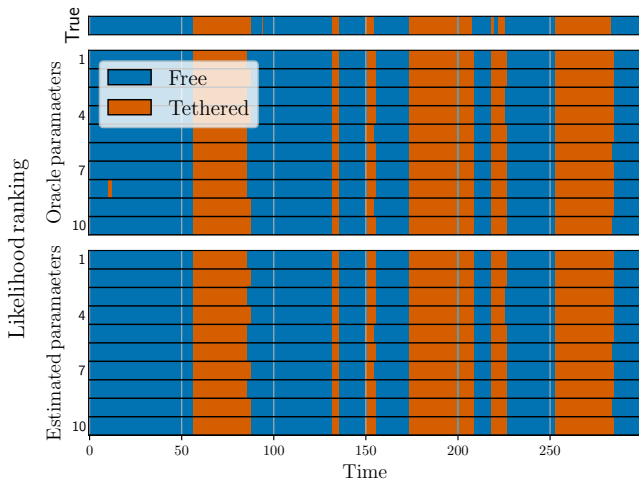

FIG. 7. The hidden paths of a single trajectory realization with $N = 300$ steps, $\tau_0 = \tau_1 = 100$, $D = 1$ and $A = 1$ and a time interval of $\Delta t = 10$. The true hidden path is at the top. In the middle are the 10 most likely hidden paths given the "oracle parameters" (see text), with the upper row being the most likely. At the bottom are the 10 most likely hidden paths given the algorithm's estimated model parameters, again with the upper row being the most likely. Each row in the middle panel differs from all the rows above it by at least one step, and the same is true for the bottom panel. Blue intervals (left and rightmost) correspond to the free state, and orange intervals correspond to the tethered state.

V. DISCUSSION AND OUTLOOK

In this work we develop and test an efficient algorithm to analyze trajectories of diffusing particles with transient tethering, identify tethering/untethering events, and estimate the physical parameters of the system. The crux of our method is to use the Viterbi algorithm to find the most likely sequence of hidden states and estimate the model parameters only for that sequence. Our algorithm successfully recovers the model parameters and is insensitive to the initial condition. It is applicable when the time interval between frames Δt is significantly shorter than the typical tethering/untethering times, and longer than, or comparable to, the equilibration time of the particle with the tethering potential.

While the spatial parameters D and A are correctly estimated to high accuracy, the temporal parameters τ_0, τ_1 are consistently overestimated as seen in Table I and Fig. 6, with an increasing overestimation as τ_0, τ_1 decrease compared to Δt. This overestimation is due to a systematic misidentification of brief tethered intervals. An example of this phenomenon is illustrated in Fig. 7 where a brief tethered interval in the true hidden path is missing from the most likely path. Due to the exponential distribution of the free and tethered interval durations in our model, short intervals are likely to occur.

It is worth noting that the overestimation is consistent between runs, as is evident in the histograms of $\hat{\tau}_0$ and $\hat{\tau}_1$, cf. Fig. 6 and Table I. One possible approach for calibrating these estimates is to use a parametric bootstrap procedure [39]. The idea is as follows: take the estimated model parameters $\hat{\Theta}$ and generate a new trajectory based on them, then re-estimate the model parameters of the new trajectory using the alternating minimization algorithm and measure the relative discrepancy of $\hat{\tau}_0, \hat{\tau}_1$. By repeating this many times, one obtains a distribution of relative errors between τ_0, τ_1 and the estimated $\hat{\tau}_0, \hat{\tau}_1$ which can then be used both for calibrating the estimates and for constructing confidence intervals.

Our method can be readily generalized to other dimensions ($d \neq 2$), by adjusting the prefactors in Eq. (9). In higher dimensions, the aforementioned misidentification of brief tethered intervals should be less prominent, since a sequence of short steps becomes less likely (no recurrence). Furthermore, our method can be generalized for other diffusive motions or different tethering potentials, like incorporating a drift term or an observational error term for position samples, as in Bernstein and Fricks's work [22].

[1] A. Einstein, Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen, Annalen der Physik **322**, 549 (1905).

[2] M. von Smoluchowski, Zur kinetischen theorie der brownschen molekularbewegung und der suspensionen, Annalen der Physik **326**, 756 (1906).

[3] J. Perrin, Mouvement brownien et réalité moléculaire, in *Annales de Chimie et de Physique*, Vol. 18 (1909) pp. 1–114.

[4] E. Kappler, Versuche zur messung der avogadro-loschmidtschen zahl aus der brownschen bewegung einer drehwaage, Annalen der Physik **403**, 233 (1931).

[5] B. Wang, S. M. Anthony, S. C. Bae, and S. Granick, Anomalous yet Brownian, Proc. Natl. Acad. Sci. USA **106**, 15160 (2009).

[6] B. Wang, J. Kuo, S. C. Bae, and S. Granick, When Brownian diffusion is not Gaussian, Nature Materials **11**, 481 (2012).

[7] A. V. Chechkin, F. Seno, R. Metzler, and I. M. Sokolov, Brownian yet non-Gaussian diffusion: From superstatistics to subordination of diffusing diffusivities, Phys. Rev. X **7**, 021002 (2017).

[8] J. Guan, B. Wang, and S. Granick, Even hard-sphere colloidal suspensions display Fickian yet non-Gaussian diffusion, ACS Nano **8**, 3331 (2014).

[9] I. Chakraborty, G. Rahamim, R. Avinery, Y. Roichman, and R. Beck, Nanoparticle mobility over a surface as a probe for weak transient disordered peptide–peptide interactions, Nano Letters **19**, 6524 (2019).

[10] I. Chakraborty and Y. Roichman, Disorder-induced Fickian, yet non-Gaussian diffusion in heterogeneous media, Phys. Rev. Res. **2**, 022020 (2020).

[11] R. Pastore, A. Ciarlo, G. Pesce, A. Sasso, and F. Greco, A model-system of Fickian yet non-Gaussian diffusion: light patterns in place of complex matter, Soft Matter **18**, 351 (2022).

[12] F. Rusciano, R. Pastore, and F. Greco, Fickian non-Gaussian diffusion in glass-forming liquids, Phys. Rev. Lett. **128**, 168001 (2022).

[13] A. Ciarlo, R. Pastore, F. Greco, A. Sasso, and G. Pesce, Fickian yet non-Gaussian diffusion of a quasi-2D colloidal system in an optical speckle field: experiment and simulations, Scientific Reports **13**, 7408 (2023).

[14] C. Yoshina-Ishii, Y.-H. M. Chan, J. M. Johnson, L. A. Kung, P. Lenz, and S. G. Boxer, Diffusive dynamics of vesicles tethered to a fluid supported bilayer by single-particle tracking, Langmuir **22**, 5682 (2006).

[15] R. Das, C. W. Cairo, and D. Coombs, A hidden markov model for single particle tracks quantifies dynamic interactions between lfa-1 and the actin cytoskeleton, PLOS Computational Biology **5**, 1 (2009).

[16] M. Ott, Y. Shai, and G. Haran, Single-particle tracking reveals switching of the HIV fusion peptide between two diffusive modes in membranes, The Journal of Physical Chemistry B **117**, 13308 (2013).

[17] S. Thapa, M. A. Lomholt, J. Krog, A. G. Cherstvy, and R. Metzler, Bayesian analysis of single-particle tracking data using the nested-sampling algorithm: maximum-likelihood model selection applied to stochastic-diffusivity data, Physical Chemistry Chemical Physics **20**, 29018 (2018).

[18] N. Granik, L. E. Weiss, E. Nehme, M. Levin, M. Chein, E. Perlson, Y. Roichman, and Y. Shechtman, Single-particle diffusion characterization by deep learning, Biophysical Journal **117**, 185 (2019).

[19] R. C. Falcao and D. Coombs, Diffusion analysis of single particle trajectories in a bayesian nonparametrics framework, Physical Biology **17**, 025001 (2020).

[20] Q. Xu, L. Feng, R. Sha, N. C. Seeman, and P. M. Chaikin, Subdiffusion of a sticky particle on a surface, Phys. Rev. Lett. **106**, 228102 (2011).

[21] M. J. Skaug, J. Mabry, and D. K. Schwartz, Intermittent molecular hopping at the solid-liquid interface, Phys. Rev. Lett. **110**, 256101 (2013).

[22] J. Bernstein and J. Fricks, Analysis of single particle diffusion with transient binding using particle filtering, Journal of Theoretical Biology **401**, 109 (2016).

[23] I.-H. Wang, C. J. Burckhardt, A. Yakimovich, and U. F. Greber, Imaging, tracking and computational analyses of virus entry and egress with the cytoskeleton, Viruses **10**, 166 (2018).

[24] P. J. Slator and N. J. Burroughs, A hidden markov model for detecting confinement in single-particle tracking trajectories, Biophysical Journal **115**, 1741 (2018).

[25] A. Callegari, C. Sieben, A. Benke, D. M. Suter, B. Fierz, D. Mazza, and S. Manley, Single-molecule dynamics and genome-wide transcriptomics reveal that NF-kB (p65)-DNA binding times can be decoupled from transcriptional activation, PLoS genetics **15**, e1007891 (2019).

[26] P. Kowalek, H. Loch-Olszewska, and J. Szwabiński, Classification of diffusion modes in single-particle tracking data: Feature-based versus deep-learning approach, Phys. Rev. E **100**, 032410 (2019).

[27] M. Jensen, Y.-Y. Wang., S. K. Lai, M. G. Forest, and S. A. McKinley, Antibody-mediated immobilization of virions in mucus, Bulletin of Mathematical Biology **81**, 4069 (2019).

[28] M. K. Huseyin and R. J. Klose, Live-cell single particle tracking of PRC1 reveals a highly dynamic system with low target site occupancy, Nature Communications **12**, 887 (2021).

[29] F. Simon, J.-Y. Tinevez, and S. van Teeffelen, Ex-Track characterizes transition kinetics and diffusion in noisy single-particle tracks, Journal of Cell Biology **222**, e202208059 (2023).

[30] A. Federbush, A. Moscovich, and Y. Bar-Sinai, Github repository, https://github.com/ybs-lab/diffusion-with-tethering (2023).

[31] L. E. Reichl, A modern course in statistical physics (Wiley-VCH GmbH, 2016) Chap. 7.3, 4th ed.

[32] Q. Li, Z. Zhu, and G. Tang, Alternating minimizations converge to second-order optimal solutions, in *Proceedings of the 36th International Conference on Machine Learning*, Vol. 97 (2019) pp. 3935–3943.

[33] A. P. Dempster, N. M. Laird, and D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, Journal of the Royal Statistical Society. Series B (Methodological) **39**, 1 (1977).

[34] D. J. C. MacKay, *Information Theory, Inference and Learning Algorithms* (Cambridge University Press, 2006).

[35] A. Viterbi, Error bounds for convolutional codes and an asymptotically optimum decoding algorithm, IEEE Transactions on Information Theory **13**, 260 (1967).

[36] N. Seshadri and C.-E. Sundberg, List Viterbi decoding algorithms with applications, IEEE Transactions on Communications **42**, 313 (1994).

[37] M. Roder and R. Hamzaoui, Fast tree-trellis list Viterbi decoding, IEEE Transactions on Communications **54**, 453 (2006).

[38] D. Golod, *The k-best paths in Hidden Markov Models. Algorithms and Applications to Transmembrane Protein Topology Recognition.*, Master's thesis (2009).

[39] B. Efron and R. Tibshirani, *An Introduction to the Bootstrap* (Chapman & Hall/CRC, Philadelphia, PA, 1994).