



HAL
open science

Variational inference of fractional Brownian motion with linear computational complexity

Hippolyte Verdier, François Laurent, Alhassan Cassé, Christian L Vestergaard,
Jean-Baptiste Masson

► To cite this version:

Hippolyte Verdier, François Laurent, Alhassan Cassé, Christian L Vestergaard, Jean-Baptiste Masson. Variational inference of fractional Brownian motion with linear computational complexity. *Physical Review E*, 2022, 106, pp.055311. <10.1103/PhysRevE.106.055311>. <hal-03612918v3>

HAL Id: hal-03612918

<https://hal.science/hal-03612918v3>

Submitted on 17 Feb 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Variational inference of fractional Brownian motion with linear computational complexity

Hippolyte Verdier*

*Decision and Bayesian Computation, USR 3756 (C3BI/DBC) & Neuroscience department
CNRS UMR 3751, Institut Pasteur, Université de Paris, CNRS, Paris, France and
Histopathology and Bio-Imaging Group, Sanofi, R&D, Vitry-Sur-Seine, France*

François laurent, Christian L. Vestergaard, and Jean-Baptiste Masson†

*Decision and Bayesian Computation, USR 3756 (C3BI/DBC) & Neuroscience department
CNRS UMR 3751, Institut Pasteur, Université de Paris, CNRS, Paris, France*

Alhassan Cassé

Histopathology and Bio-Imaging Group, Sanofi, R&D, Vitry-Sur-Seine, France

We introduce a simulation-based, amortised Bayesian inference scheme to infer the parameters of random walks. Our approach learns the posterior distribution of the walks' parameters with a likelihood-free method. In the first step a graph neural network is trained on simulated data to learn optimised low-dimensional summary statistics of the random walk. In the second step an invertible neural network generates the posterior distribution of the parameters from the learnt summary statistics using variational inference. We apply our method to infer the parameters of the fractional Brownian motion model from single trajectories. The computational complexity of the amortised inference procedure scales linearly with trajectory length, and its precision scales similarly to the Cramér-Rao bound over a wide range of lengths. The approach is robust to positional noise, and generalises to trajectories longer than those seen during training. Finally, we adapt this scheme to show that a finite decorrelation time in the environment can furthermore be inferred from individual trajectories.

I. INTRODUCTION

Fractional Brownian motion (fBm) [15, 48] is a paradigmatic model of anomalous transport. It is a non-Markovian Gaussian process characterised by stationary increments and long temporal correlations in the noise driving the process. It allows capturing long-range temporal correlations in the dynamics of a walker or its environment, and it is a model of choice to describe a multitude of dynamic processes in numerous scientific fields [4, 10, 11, 17, 19, 24, 25, 29–31, 33, 41, 44, 56, 64, 71, 72, 75, 79]. Following the classification given in [51] of the three main sources of anomalous diffusion, the anomalous dynamics of fBm stems from the statistical dependency of the displacements at all time scales. Since fBm is a Gaussian process, it admits an analytical expression of the joint likelihood of the recorded signal. It is thus an ideal model to investigate the performance of approximate schemes to infer anomalous diffusion, such as variational inference or machine-learning-based approaches, since it allows direct comparison to statistically optimal exact inference.

The position of a random walker undergoing fBm is described by a Langevin equation [15] of the form

$$\frac{dr(t)}{dt} = \sqrt{K_\alpha} \eta(t), \quad (1)$$

where η is a zero-mean Gaussian noise process with covariance $\langle \eta(t_1) \eta(t_2) \rangle = \alpha(\alpha - 1) |t_1 - t_2|^{\alpha-2}$ and K_α is a generalised diffusion constant that sets the scale of the process. fBm is self-similar and ergodic [12, 21]. However, it has been shown to exhibit transient non-ergodic behaviour when confined [12, 32] and it is worth noting that the ergodic regime is witnessed only after a long transient passage exhibiting non-ergodic properties [12, 27, 42, 45]. The noise η is negatively correlated in the subdiffusion regime ($\alpha < 1$), while it is positively correlated in the super-diffusion regime ($\alpha > 1$).

Methods for estimating a random walk's parameters can roughly be divided into two types: heuristic approaches using features extracted from the trajectories [30, 37, 50, 54, 65], and likelihood-based (e.g., Bayesian) approaches [36, 40, 47, 66]. Each has its strengths and weaknesses. Likelihood-based approaches are provably asymptotically optimal, but they are often computationally intensive and are only applicable to random walk models that have a tractable likelihood. Feature-based approaches are typically computationally cheaper, and they can be applied to a much larger range of models since they do not require a tractable likelihood. However, they are generally not statistically efficient, are prone to bias when used on experimental data and their precision can be difficult to evaluate. It is worth noting the rapid progress of machine learning based approaches [53, 54], which fall in the category of feature-based approaches, and which allow to learn high quality features to perform both parameter estimation and model classification. While such machine learning ap-

* hverdier@pasteur.fr

† jbmasson@pasteur.fr

proaches generally outperform handcrafted features on numerically generated data, it remains difficult to evaluate their actual performance and robustness on empirical data.

Here, we develop an amortised variational inference approach to estimate the parameters of a fBm from a single recorded trajectory. The objective of this paper is three-fold : i) we explored the possibility, via amortisation, to reduce the marginal computational cost of inference from quadratic to linear in the trajectory length, ii) we compared the performance of a such estimator with the Cramer-Rao bound and we assessed information retrieval by a linear computational scheme and iii) we investigated the possibility of using this variational approach to infer the posterior distribution of a finite decorrelation time of the walker’s dynamics.

The inference scheme relies on a graph neural network (GNN) trained on simulations of fBm realisations, which encodes a vector of summary features for each trajectory. The encoder’s architecture allows it to naturally capture long-range interactions while retaining a linear scaling of the computational complexity with the length of the trajectories. It is associated with an invertible network, which generates the posterior distribution from the summary features using a variational objective. We applied the method to the fBm model because the tractability of its likelihood allows us to compute the Cramér-Rao bound which provides a lower bound on the variance of unbiased estimators. We show that the approach attains near-optimal performance as compared to exact likelihood-based inference and to the Cramér-Rao bound. We furthermore discuss the latent space structure learned by the summary network and its ability to encode physical properties. We test the applicability of the approach to trajectories corrupted by positional noise and its potential to generalise to trajectories that are longer than those seen during training. Finally, we extend the inference procedure to capture a finite decorrelation time in the dynamics which may typically arise in physical environments.

II. AMORTISED BAYESIAN INFERENCE OF FBM

A. fBM : a tractable random walk model

In the context of parameter estimation, Bayesian inference uses Bayes’ theorem to compute the posterior probability distribution of the parameters θ given recorded data \mathbf{R} (here a trajectory) and a probabilistic model of these data,

$$p(\theta|\mathbf{R}) = \frac{p(\mathbf{R}|\theta)p(\theta)}{p(\mathbf{R})}. \quad (2)$$

Equation equation (2) relates the posterior distribution, $p(\theta|\mathbf{R})$ to the likelihood $p(\mathbf{R}|\theta)$, the prior $p(\theta)$ and the evidence $p(\mathbf{R})$. Here, we only consider one single model,

i.e., the fBm, and thus do not explicitly refer to it. The principle of amortised inference [14] is to split the estimation of the posterior $p(\theta|\mathbf{R})$ into two independent steps. The first is computationally costly and involves learning an approximate posterior density $\hat{p}(\theta|\mathbf{R})$ from numerically generated data. Then, the second step consists in running the pre-trained approximate system on the experimental data to infer the posterior density, assuming that they are similar to the training data.

A tractable likelihood can be computed for fBm: considering a trajectory $\mathbf{R} = (r_0, r_2, \dots, r_N)$ to be a 1-dimensional time-series of positions r_i recorded at equidistant points in time $t_i \in \{0, \Delta t, 2\Delta t, \dots, N\Delta t\}$, the likelihood of a trajectory reads

$$p(\mathbf{R}|\theta) = \frac{1}{(2\pi)^{N/2} \sqrt{\det \Sigma(\theta)}} \exp\left(-\frac{1}{2} (\Delta\mathbf{r})^\top \Sigma(\theta)^{-1} \Delta\mathbf{r}\right), \quad (3)$$

where $\Delta\mathbf{r} = (\Delta r_1, \dots, \Delta r_N)^\top$, with $\Delta r_i = r_i - r_{i-1}$ the individual displacements. Then, $\theta = (K_\alpha, \alpha)$ are the fBm’s parameters to infer, and Σ is the displacements’ covariance matrix whose coefficients are given by

$$[\Sigma(\theta)]_{ij} = K_\alpha \Delta t^\alpha (|i - j + 1|^\alpha + |i - j - 1|^\alpha - 2|i - j|^\alpha). \quad (4)$$

B. Amortised inference

We choose to rely on a likelihood-free approach to amortise our inference procedure. This may seem a counter-intuitive choice for the precise case of fBm because the likelihood is analytically tractable – other inference methods leverage this specificity, see for example [67] – but our method has the advantage of relying solely on computations of linear complexity. Furthermore, the approach is also directly portable to more complex problems for which a tractable likelihood may not be available or may be computationally too costly. Indeed, likelihood-free inference is a method of choice to address such problems. As more and more complex models are encountered in numerous fields of science, the field of simulation-based inference [14] is growing very rapidly to address the associated challenging inverse problems. The shift towards amortisation of the likelihood is notably driven by new tools and conceptual approaches derived from machine learning [1, 57]. Here, applying the method to a problem with a tractable likelihood allows us to compare its performance to the optimal one, derived using the Cramér-Rao bound.

The architecture of the amortised variational inference scheme, allowing to estimate the posterior distribution, is shown in Figure 1. It is based on the recently introduced Bayes Flow (BF) [61] procedure. In this framework, a first neural network, the “summary network” (working as an encoder), computes a vector of summary statistics from observations. Here, the summary network is a GNN (Fig. 1A) taking as input a trajectory and outputting a summary statistics vector \mathbf{h} whose dimension




FIG. 1. **Model Architecture.** A: Construction of a graph from a single trajectory (left). Positions, colored according to time, are treated as nodes for which features \mathbf{X} are computed (middle). Nodes are then connected by edges (grey lines) following a given wiring scheme, and edge features \mathbf{Y} are computed. Edges terminating at the trajectory’s last point are shown in red. Feature matrices for both nodes (\mathbf{X}) and edges (\mathbf{Y}) are depicted in small insets with color coded values. B: Summary network extracting information from the trajectory’s graph. It consists in several graph convolution layers (in purple), a pooling layer (green) and a multi-layer perceptron (orange). The vector of statistics it outputs is indicated by \mathbf{h} . C: General structure of the model, with the summary network (light blue) extracting from the trajectory $\mathbf{r}(t)$ the summary statistics vector \mathbf{h} , which in turn parameterises the invertible network. During the training phase, the invertible network is used from left to right (i.e., from the parameter’s manifold to an easily sampled one), and in inference mode it is used from right to left.

is independent of the trajectory length. This vector then parameterises an invertible transformation between an easily sampled distribution (multivariate normal) and the posterior distribution of the parameters (Fig. 1B). The full procedure generates a posterior distribution of the parameters. Such flow-based approaches, derived from normalising flows [35], have the advantage of providing an estimation of the posterior without requiring extensive sampling. The whole neural network is trained on numerically generated data and can then be used for inference. In the two following subsections, we first present the GNN (Section II C) and then the parameterisable invertible network (Section II D). The parameters of both these parts of the neural network are set during an upfront

training phase, detailed below, during which fBM trajectories having the same parameters as those on which we seek to later perform inference are presented to the network.

C. Graph neural network for learning summary statistics

GNNs have been introduced to model and analyse graphs, meshes and point clouds [26, 34, 59]. They are well suited to capture geometric properties from point clouds and other datasets of variable size [13, 59], they can keep a sparse architecture while encoding long tem-

poral correlations [5, 46] and they exhibit good performance with a limited number of parameters compared to other modern architectures. In the following, we present the specific GNN architecture used for the summary network in this work. It is an updated version of the GRATIN architecture (as "Graphs on Trajectories for Inference") which we introduced in [69]. We summarise in 1c the differences between this version of the GRATIN architecture and that of the encoder of [69]. We refer to this previous work and references therein the reader interested by more thorough explanations about graph convolutions, which are the core operations performed by GNNs. Throughout this paper, we use the GRATIN acronym to refer both to the architecture of the summary statistics network and to the inference scheme as a whole.

As indicated by their name, GNNs process graphs, and the first step of our inference is thus to build a graph from a trajectory. To do so, we represent each trajectory $\mathbf{R} = (r_0, r_1, r_2, \dots, r_N)$ by a directed graph $G = (V, E, \mathbf{X}, \mathbf{Y})$. Here $V = \{1, 2, \dots, N\}$ is the set of nodes, each corresponding to a recorded position of the observed walker. $E \subseteq \{(i, j) | (i, j) \in V^2\}$ is the set of edges connecting pairs of nodes.

A vector of features, $\mathbf{x}_i^{(0)}$ (of size n_x), is first attached to each node, encapsulating information relating to the i -th position of the trajectory and to the displacements that led to it. There are six such node features (i.e. $n_x = 6$): the normalised time i/N , the maximal jump size since the beginning of the trajectory \mathbf{R} (normalised by $s(\mathbf{R})$, the standard deviation of jump sizes), the distance to origin and the maximal distance to origin up to point i , both normalised by $s(\mathbf{R})\sqrt{i}$, which is proportional to the square root of the expected square displacement of a Brownian walker over this time span. These are complemented with three features indicative of the moments of the distribution of step sizes observed up to node i , detailed in supplementary 1a. Thus, $\mathbf{X}^{(0)} = (\mathbf{x}_1^{(0)}, \mathbf{x}_2^{(0)}, \dots, \mathbf{x}_N^{(0)})$ is the (N, n_x) matrix of initial node feature vectors.

Similarly, $\mathbf{Y} = (\mathbf{y}_1^{(0)}, \mathbf{y}_2^{(0)}, \dots, \mathbf{y}_{|E|}^{(0)})$ is a matrix of edge features, $\mathbf{y}_e^{(0)}$, each associated to an edge e in E . The features vector of a given edge $e = (i, j)$, $\mathbf{y}_e^{(0)}$, of size $n_y = 6$, encapsulates information about the trajectory's course between the two nodes i and j it connects. More precisely, edge features are the time difference $j - i$, the distance between its extremities normalised by $s(\mathbf{R})\sqrt{j - i}$, the correlation of jumps normalised by $s(\mathbf{R})^2$, as well as three other features related to the moments of the distribution of step sizes observed between nodes i and j .

The edges in G are chosen such that incoming edges of each node originate only from nodes in the past (i.e., respecting causality): node i receives connections from nodes $i - \lfloor i^{k/(d-1)} \rfloor$, for $k \in \{0, 1, \dots, d-1\}$ where $\lfloor \cdot \rfloor$ denotes the integer part and d is the maximum incoming degree parameter: each node receives a maximum of d edges, and they are thus less than dN edges in the graph

(in the following, we chose $d = 20$). This scale-free wiring of nodes allows the network to capture long range correlations, and provides a good control on the way the networks structures the data. The graph structure furthermore allows to summarize the geometry of the trajectory without specifying coordinates, by resorting to distances (carried by the edges). Thus, the graph representation of trajectories (with nodes and edge features) taken as input by the GRATIN summary network, is rotation-invariant. This is an important advantage given that the inference problem indeed has a such invariance – respecting symmetries of the data does in general help neural networks to converge [28, 43]

While the training is specific to the dimension of the random walk, the GNN architecture can keep the same features initialisation and number of parameters.

A key point about the graph construction procedure (wiring and features initialisation) is that its algorithmic complexity is linear in the number of nodes. This is required by the objective of providing an estimator of linear complexity.

Following the graph initialisation step, the summary network performs several graph convolution operations [20, 73, 78]. It then passes the learnt node feature vectors as inputs to a pooling layer that aggregates features across all nodes of a trajectory graph into a fixed-length vector. The vector is finally passed through a multi-layer perceptron to obtain the summary statistics vector $\mathbf{h} = g_\psi(\mathbf{R})$, where ψ denotes the neural network coefficients, set during training. In section III, we show inference results done on trajectories longer than those seen by the networks during their training. While the architecture is capable to handle trajectories of any size, it will not produce reliable results on trajectories which have a such major difference with those seen during the training phase. Hence, we segment long trajectories in segments of length $N \leq 1\,000$ and average the vectors of summary statistics computed from these segments to obtain a summary statistics vector describing the complete trajectory.

D. Invertible network for generating a variational posterior density

The Bayes Flow approach relies on an invertible transformation, $f_\phi(\cdot; \mathbf{h})$, between the parameter space (in \mathbb{R}^D , with $D \geq 2$) and the prior space (in \mathbb{R}^D), on which a D -dimensional standard Gaussian density is assumed. This allows the model to output a full posterior distribution over the space of parameters, while most machine learning-based techniques aimed at analyzing random walks usually restrict to point estimates. The transformation $f_\phi(\cdot; \mathbf{h})$ is parameterised by a conditional invertible neural network (cINN) [2] made of a succession of affine coupling blocks [22] (multiple blocks sequentially applied) and maps θ to the prior conditioned on \mathbf{h} , the summary statistics of the trajectory.




FIG. 2. **Model performance.** A: Evolution of the posterior density of α inferred by the model (plain blue lines) versus the true posterior (dotted green lines) from two example trajectories with $\alpha = 0.7$ (top) and $\alpha = 1.3$ (bottom), respectively. The length of the portion of the trajectory used for inference is increased ten-fold between each panel, i.e., from left to right: $N = 10, 100, 1\,000$. B: Density of $(\alpha, \mathbb{E}[\hat{\alpha}])$ for 10 000 trajectories of different lengths. From left to right, $N = 10, 50, 250, 1\,000$. C: histograms of $\mathbb{E}[\hat{\alpha}]$ for several values of α (from left to right, $\alpha = 0.5, 0.8, 1.2, 1.5$, marked by the red vertical lines), with trajectories of different lengths ($N = 10$ (blue), 50 (orange), 250 (green), 1 000 (red)). Histograms are representative of 10 000 trajectories each. D: Comparison of the MSE of our estimator of α with the optimal variance given by the Cramer-Rao bound, over the range of values taken by alpha. For each of 30 values of alpha, the MSE is an average over 10 000 trajectories.) E: Blue curves: evolution, with the length N of trajectories, of the mean square error (MSE) of the mean posterior $\mathbb{E}[\hat{\alpha}]$ obtained with several architectures of summary networks. for each value of N , 100 000 trajectories are used to compute the averages. Green curves: Cramér-Rao bound for an unbiased estimator of α , and MSE of the α of maximum-likelihood (on an evenly spaced grid of 200 values between 0.1 and 1.9).

By design, these blocks can be inverted and the determinant of the Jacobian matrix \mathbf{J}_{f_ϕ} of the transformation retrieved from the forward pass. During training we seek to approximate the true posterior $p(\theta|\mathbf{R})$ by the learnt posterior $p_\phi(\theta|\mathbf{R}) = \exp\left(-\frac{\|f_\phi(\theta;\mathbf{h})\|_2^2}{2}\right)$. The loss function is the Kullback-Leibler divergence between $p(\theta|\mathbf{R})$ and $p_\phi(\theta|\mathbf{R})$ which reads as

$$\mathcal{L}(\mathbf{R}) = \frac{1}{2}\|f_\phi(\theta;\mathbf{h})\|_2^2 - \log|\det\mathbf{J}_{f_\phi}|, \quad (5)$$

where $\mathbf{h} = g_\psi(\mathbf{R})$ and $\|\cdot\|_2$ denotes the Euclidean norm. Sampling the posterior distribution first requires computing \mathbf{h} from the trajectory \mathbf{R} , and then generating the required number of sample as $\theta = f_\phi^{-1}(\mathbf{z};\mathbf{h})$ with \mathbf{z} drawn from a standard D -dimensional Gaussian distribution.

Details about the implementation are presented in Supplementary Material 1., and the code is available online : <https://gitlab.pasteur.fr/hverdier/gratin-fbm>

III. ESTIMATION OF THE ANOMALOUS EXPONENT

We evaluate the performance of our variational inference procedure on numerically generated trajectories. Estimating the anomalous exponent α is the most challenging part of the inference, and we thus focus on this here, although our approach infers a joint posterior density for $\theta = (K_\alpha, \alpha)$. Trajectories used for training and evaluation are generated using the Davies-Harte method [18] implemented in the Python `fbm` package with $\alpha \sim \mathcal{U}(0.1, 1.9)$ and $\log_{10}(K_\alpha) \sim \mathcal{U}(-2, 2)$.

Figure 2A shows the inferred posteriors of α on portions of increasing length of two example trajectories. The amortised posterior is consistent with the exact one, and both become increasingly peaked around the true value of α as the length of the trajectory increases. The inferred posterior distributions do not exhibit broad tails or divergences, and are proper distributions, i.e., they are normalisable.

Using the likelihood shown in 3 we computed the Cramér-Rao bound of the inference problem, which gives a lower bound of the variance achievable by an unbiased estimator of the parameters θ (see Supplementary Material 3). We computed $b_N(\alpha)$ the lower bound of the variance on α for estimators processing fBM realisations of length N with an anomalous diffusion exponent α . Besides, reducing for each trajectory the inferred posterior distribution to its mean, $\mathbb{E}[\hat{\alpha}]$, we compute the mean square error (MSE) $\langle(\mathbb{E}[\hat{\alpha}] - \alpha)^2\rangle$ of the inference. On Fig. 2D, we compare this quantity (for $N = 100$) to the MSE of our amortized inference based on the GRATIN summary network. We observe that the optimal variance critically depends on α and that our inference is close from optimal on trajectories with $\alpha \approx 1$.

In Fig. 2E, we show the mean MSE of our inference on trajectories with lengths varying across two orders of

magnitudes, with α uniformly sampled in the interval (0.1, 1.9). We observe that the MSE follows a power-law decrease and is close to the mean optimal variance (obtained by averaging, for each N the values of $b_N(\alpha)$ over the range of values of α). The amortised inference is suboptimal (as expected from any variational inference), but its variance shows a fast decreasing trend similar to the Cramér-Rao bound, i.e., close to $\propto 1/N$. We compared several architectures of summary networks : GRATIN, which is described above, GRATIN-NF (as "no features"), which is a version of GRATIN with a minimal set of features (see supplementary material 1 a) and an architecture based on 1D convolutions and LSTM layers ("Long Short Term Memory", [8]), inspired from [3] and [55], who proposed similar architectures to study anomalous diffusion (see details in supplementary material ??). We observe that features do improve the quality of the inference performed by GRATIN. The LSTM summary network provides slightly more precise predictions than the GRATIN one, but this is at the cost of a lesser interpretability of the inner computation performed by the network and a greater number of parameters.

An advantage of a BayesFlow-based framework is that, by averaging summary statistics vectors of segments of long trajectories, the network is able to extract information about longer trajectories than those seen during training. We show in figure 3A that the spread of mean inferred predictions keeps shrinking as N increases. We quantify this in figure 3B, where the evolution of the mutual information (MI) between the true and the mean inferred α is shown as a function of trajectory length. To estimate MI, we used the procedure described in [39] implemented in [58]. Nevertheless, as segments of a long trajectories are correlated, so are their summary statistics vectors. Hence, their average does not necessarily lay in the same region of the latent space as representations of segments of same size taken from individual trajectories with the same value of α . This results in a slight systematic bias of the prediction, which we corrected using a simple polynomial.

The rotation invariance which the GNN enforces is especially relevant on trajectories of higher dimension. To confirm the interest of a such constraint, we have trained inference models based on GRATIN and LSTM summary networks on trajectories of increasing dimension (and fixed length $N = 100$) and observed that models based on GRATIN yielded better results than LSTMs for trajectories of dimension greater than three, as shown in figure 3D. The LSTM architecture seems to encounter convergence issue when processing trajectories of high dimensions. We kept the same hyperparameters of training throughout all experiments, i.e. a batch size of 2 048 and a learning rate of 0.01, which is reduced by a factor 0.2 when the loss plateaus for 15 epochs. Training was stopped after 30 epochs of 100,000 trajectories. While a thorough hyperparameter optimization could reveal the configuration which enables the LSTM architecture to properly converge, this highlights the robustness of the

GNN architecture. Note that interest in GNN is based on physical interpretability of the latent space, linear complexity with trajectory length and the possibility to implement physically relevant symmetries within the Graph structure. There is little doubt that numerous machine learning architecture may be found to be more efficient in specific point prediction processes.

The learnt summary statistics \mathbf{h} , used to compute the posterior distribution, constitutes a low-dimensional representation of a trajectory. The structure of the manifold on which summary statistics vector lie is revealing of the way the encoder represents information extracted from the trajectories. An assumption in representation learning [6] is that interpretable representation lead to better generalisation. We projected \mathbf{h} onto a 2D plane using UMAP [49] (a non-linear dimensionality reduction algorithm) and mapped α on it (see Supplementary Fig. S1). We see that the latent space is organised according the value of α , a good indication that the learning process properly captured the underlying physical properties. We tested the robustness of the inference procedure when applied to trajectories corrupted by positional noise. We show in Supplementary Fig. S2 the evolution of the MSE of the amortised inference of α and compare it with the corresponding Cramér-Rao bound. The precision of the amortised inference procedure closely follows the lower bound set by the Cramér-Rao inequality. This was obtained by training models specifically on trajectories corrupted with increasing amounts of noise.

An important attribute of our approach is that, passed the amortisation (training) step, the computational complexity of inferring the parameters of a novel trajectory is linear in terms of the length of this trajectory. To show this, we subdivide the inference procedure into three steps: (i) initial feature evaluation, (ii) forward pass through graph convolutions and pooling, and (iii) operations on summary statistics to generate the posterior. (i) The initial evaluation of node and edge features requires $O(N + |E|)$ time and memory, where N is the number of nodes (for a trajectory of $N + 1$ points) and $|E|$ is the number of edges. Here $|E| \propto N$ by design (the in-degree of nodes is bounded), so this step has $O(N)$ complexity. (ii) The forward pass through the graph convolutions and the following pooling of node features requires $O(|E|)$ operations and memory slots, and hence this step also has $O(N)$ complexity. (iii) The latent space is of fixed dimensions, and hence all operations after the pooling layer have $O(1)$ complexity. The global complexity of the amortised architecture is thus linear with respect to the number of points in the trajectory.

In comparison, calculating the exact likelihood [Eq. equation (3)] requires evaluating the determinant $\det \Sigma(\boldsymbol{\theta})$ and the quadratic form $(\Delta \mathbf{r})^\top \Sigma(\boldsymbol{\theta})^{-1} \Delta \mathbf{r}$, which can be done in $O(N^2)$ time [52]. This makes exact inference prohibitively expensive for very long trajectories, where our amortised inference scheme may instead be used (Fig. 2C). Note furthermore that for many models the exact likelihood cannot be calculated at all, in which

case approximate inference is the only route possible.

IV. ESTIMATION OF A FINITE DECORRELATION TIME

When considering fBm as a model of biomolecule random walks, we have to keep in mind that many physical environments might exhibit a finite decorrelation time τ_c possibly stemming from motion occurring outside a polymer-dominated environment [76] or from changes of conformations of the biomolecule altering the nature of its interactions. The characteristic time bears information on the local environment’s physical properties, and it might be spatially dependent or specific to interactions with local partners. In practice, inferring τ_c from individual trajectories is challenging. Autocorrelation-based approaches for example give incomplete results on individual trajectories as the limited number of points prevents proper averaging [16, 63]. The power spectrum density (PSD) is a relevant quantity as well, but can only be estimated with sufficient precision for trajectories longer than those considered here by several orders of magnitude. Besides, [38] show that PSD curves computed on an ensemble of superdiffusive fBM trajectories cannot be simply averaged, hence the interest of inferring τ_c directly from single trajectories.

We adapted the amortised inference procedure presented above to infer (α, τ_c) instead of (K_α, α) . We left out K_α here since it is simply a scale factor and can be removed by rescaling the trajectories. We used the same node and edge features as above, ensuring the procedure’s linear computational complexity. A finite decorrelation time was modeled by multiplying the autocovariance of the fBm by an exponential factor, $\min(1, e^{\tau_c - \tau})$, where τ is the time difference. Examples of the autocorrelation function for several values of α and τ_c are given in Supplementary Fig. S3. The modified covariance matrix of the displacements thus reads

$$[\Sigma(\alpha, \tau_c)]_{i,j} = \min\left(1, e^{\tau_c - |i-j|\Delta t}\right) \times (|i-j+1|^\alpha + |i-j-1|^\alpha - 2|i-j|^\alpha),$$

where we have ignored the scale factor K_α . As there is no simple means to quantitatively relate the length of a trajectory to the difficulty of inferring its finite correlation time, we performed this inference on trajectories of fixed length $N = 1\,000$, with τ_c integer-valued and ranging from 5 to 50. We chose this value of N firstly because the length of the trajectory must exceed the correlation time by at least an order of magnitude in order for some information to be retrievable and secondly because the simulation of trajectories has a quadratic complexity in N , deterring one to include too long trajectories in the simulation-greedy training phase. We chose as priors $\log(\tau_c) \sim \mathcal{U}(\log(5), \log(50))$ and $\alpha \sim \mathcal{U}(0.4, 1.6)$. We restricted the range of α compared to the previous sections because the modified covariance matrices were

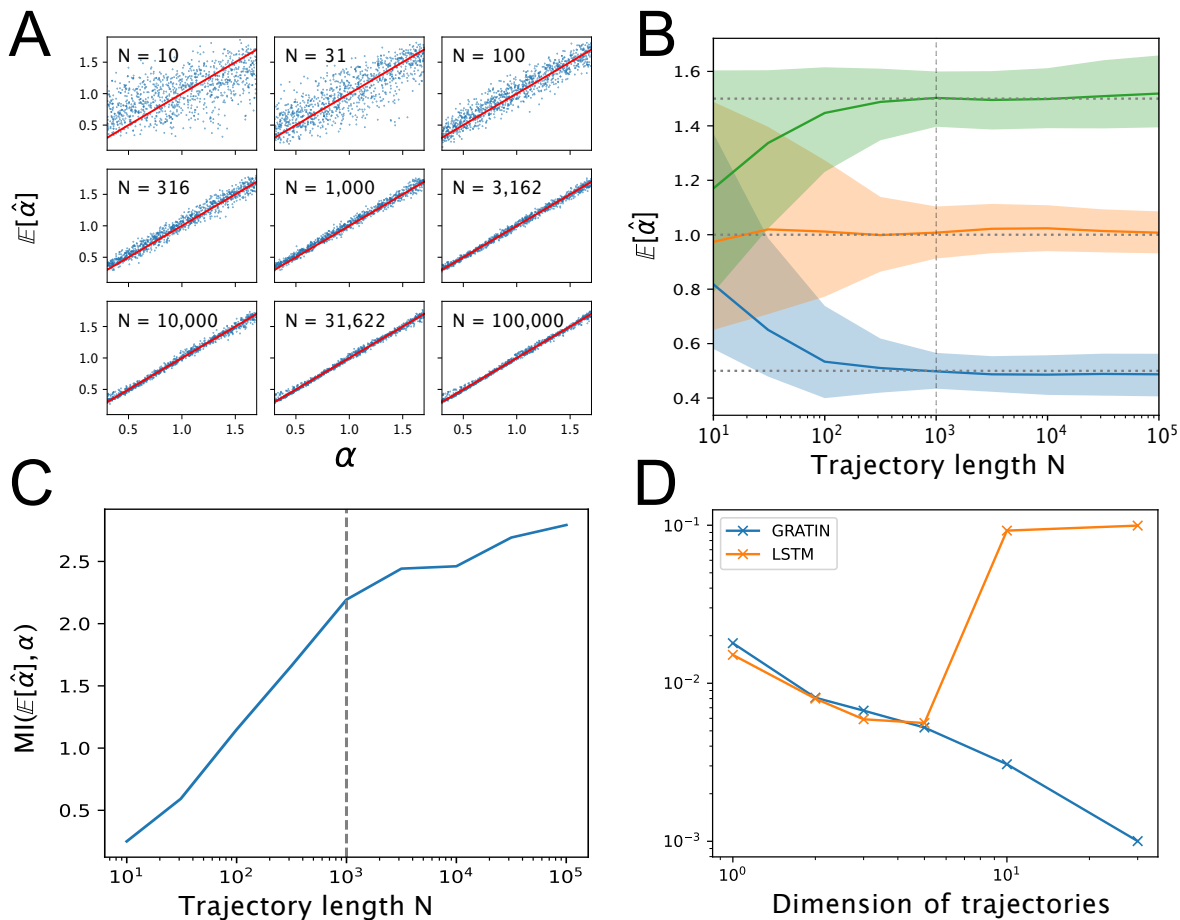


FIG. 3. A: Comparison of mean inferred values and true values of α , for trajectory lengths spanning four orders of magnitude. The red line is the diagonal $\alpha = \mathbb{E}[\hat{\alpha}]$. B: Convergence towards the true value of α as function of trajectory length. The model was trained on trajectories of length $10 \leq N \leq 1\,000$; longer trajectories were cut in segments of length $\leq 1\,000$ and their summary vectors averaged to obtain a single posterior for the whole trajectory. Shaded areas represent the 5% & 95% quantiles. C: Mutual information between α and $\mathbb{E}[\hat{\alpha}]$. D: Performance of models based on GRATIN and LSTM encoders as a function of the trajectory dimension ($D = 1, 2, 3, 5, 10, 30$).

not invertible for extreme values of α . We simulated trajectories by sampling the steps from the multivariate Gaussian distribution $\mathcal{N}(\mathbf{0}, \Sigma(\alpha, \tau_c))$ using a Cholevsky decomposition: $\Delta \mathbf{r} = L \mathbf{u}$, where L is the Cholevsky decomposition of $\Sigma(\alpha, \tau_c)$ and \mathbf{u} is a Gaussian vector drawn from the multivariate normal.

Examples of inferred joint posteriors of $(\alpha, \log(\tau_c))$ are shown in Fig. 4A. The precision of the inference of α remains high despite the limited correlations, while the order of magnitude of τ_c is often retrieved (except around $\alpha \approx 1$). The example trajectories shown on the rightmost column give a good intuition of the effect of the finite correlation time as they were generated with the same random vector \mathbf{u} (only the autocovariance matrix varied): all three trajectories have the same global shape, only fine-grained dynamics reveal the difference of anomalous exponent.

We compared our estimator with the maximum-likelihood one, which consists in choosing the value of

(α, τ_c) that maximises the likelihood of the observed trajectory. In practice, we computed the log-likelihood on a grid of values of α and τ_c to find the values yielding the maximal likelihood, with α taking 30 regularly spaced values between 0.4 and 1.6, and τ_c taking 25 geometrically-spaced values in its range. As shown in Figure 4C, our amortised inference yields a slightly more biased estimate than the maximum-likelihood estimator (when taking the mean of the posterior distribution) but has a smaller variance. When $\alpha = 1$, successive increments are completely independent of each other and there is thus no information to retrieve regarding τ_c . This is observable on the lower panel of Figure 4, both by looking at the Cramér-Rao bound, which diverges, and at the variance of our estimator, which is maximal at $\alpha = 1$.

V. DISCUSSION




FIG. 4. **Performance of the anomalous exponent and decorrelation time estimation** A: left column : joint posteriors of α and $\log_{10}(\tau_c)$, middle columns : marginal inferred posteriors (plain lines) and true ones (dashed lines), right column: corresponding trajectories. True values of the parameters are indicated in red. From top to bottom, (α, τ_c) equals $(0.6, 5)$, $(1, 20)$, $(1.4, 40)$. B: Variance and MSE of our inference of $\log(\tau_c)$ plotted as a function of α , compared with the Cramér-Rao bound for an unbiased estimator. For each of 20 values of α , 600 trajectories with $\log(\tau_c)$ sampled uniformly between $\log(5)$ and $\log(50)$ were used to compute the mean square error on the inference of $\log(\tau_c)$. C: Comparison of the values of $\log(\tau_c)$ inferred by our method (in blue) and by a maximum-likelihood (ML) estimator (in green). The thick line represents the mean across all trajectories, while the filled regions correspond to the first and last quartiles. The anomalous diffusion exponent α is drawn from a uniform distribution between 0.4 and 1.6. We considered, for each value of τ_c , 100 trajectories for the ML estimation (computationally intensive) and 6,000 for our amortized inference. The most plausible values of τ_c were chosen on a grid of 20 values of α and 20 values of τ_c .

Simulation-based inference coupled with machine learning are a promising avenue to address challenging inverse problems. When applied to intractable systems (i.e., systems for which no closed-form expression of the likelihood is known), this combination allows splitting the inference task into two steps. In the first phase, computationally intensive, simulations produce artificial data on which a neural network is trained to approximate the posterior distribution of the parameters using a variational objective. In the second step, which is computationally fast, inference is performed on experimental data by a direct forward pass through the trained network, yielding the estimated posterior distribution. The procedure is statistically efficient if the numerical data match the properties of experimental one and if the variational inference is able to capture the complex relations that might exist between the variables to be inferred.

There are two main challenges associated with amortised approaches. First, training variational inference often consists in minimising a Kullback-Leibler divergence between the approximate distribution and the real (unknown) one [7]. Optimising such a non-convex function is challenging and is not generally guaranteed to converge towards a global optimum. The second challenge is linked to interpretability. Both the models used to learn the summary statistics and the variational posterior distribution are generally intractable. There is thus limited insurance that the process does not misbehave, especially when applied to real experimental data. Evaluation of the exact posterior distribution using sampling, such as in approximate Bayesian computation, may however lead to similar problems due to the difficulty of properly sampling complex likelihood landscapes. Discrepancies between the data on which the inference is to be run and the simulated data used for training, likely to harm the inference’s relevance, could be detected using methods such as the one we introduce in [70], based on statistical tests between summary vectors.

We here used fBm to quantify the performance of our amortised inference approach. We chose to focus on fBm both due to its paradigmatic status as an anomalous random walk model and because it has a tractable likelihood, allowing us to compare our amortised method to exact likelihood-based inference and to the Cramér-Rao bound on estimator precision. We advocate more generally for the use of exactly solvable random walk models, such as the fBm, as benchmarks to evaluate the performance of machine-learning based inference methods.

We showed that our amortised inference can successfully be applied to infer the parameters of fBm, with a precision that is lower than the Cramér-Rao bound but which increases with a scaling that is similar to it. Our algorithm has a linear complexity in the length of the trajectory and can be applied to trajectories of any length at

inference time, even if the algorithm has not been specifically trained on trajectories of the same length. We furthermore showed that our amortised approach could be used to efficiently infer the parameters of a more realistic fBm-type model with a finite decorrelation time.

Our amortised inference framework can be used for any random walk model, even for models that do not have a tractable likelihood, provided that they can be simulated efficiently enough to provide a large number of trajectories for training. In all cases, our approach retains its linear computational complexity at inference time. For random walk models with intractable likelihoods, only empirical evaluation of the performance will in general be possible. Thus, it is not possible to make absolute statements about the statistical efficiency of the approach in these cases.

Beyond random walks, amortised inference can more generally be instrumental in providing posterior distributions for models of complex systems with fractional noise and/or long memory. Recent work by the authors of BayesFlow has extended its scope of application to model comparison problems [60]. Numerous challenges remain to be addressed to standardise the optimisation of the variational inference, especially in cases where some parameters are not sufficiently constrained by the data or when there are sloppy directions in the parameter space [74]. Furthermore, variational inference does not necessarily lead to physically realistic parameters. Ensuring the physics-informed [9, 62] nature of the inference may require imposing constraints on the network generating the summary statistics, but our results show that the network is able to learn physically meaningful features without inductive bias. Finally, the statistical efficiency of amortised approaches will depend on the ability of numerically generated data to match experimental observations.

Acknowledgments. We thank Thomas Blanc, Mohamed El Beheiry, Sridhar Turaga, Hugues Berry, Raphael Voituriez & Bassam Hajj for helpful discussions. This study was funded by the Institut Pasteur, *L’Agence Nationale de la Recherche* (TRamWAY, ANR-17-CE23-0016), the INCEPTION project (PIA/ANR-16-CONV-0005, OG), and the “*Investissements d’avenir*” programme under the management of Agence Nationale de la Recherche, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute).

The funding sources had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Conflicts of interest. Hippolyte Verdier and Alhassan Cassé are Sanofi employees and may hold shares and/or stock options in the company. The other author declare to have no financial or non-financial conflicts of interest.

- [1] Alsing, J., Charnock, T., Feeney, S., Wandelt, B., 2019. Fast likelihood-free cosmology with neural density estimators and active learning. *Monthly Notices of the Royal Astronomical Society* URL: <https://doi.org/10.1093/mnras/stz1960>, doi:doi:10.1093/mnras/stz1960.
- [2] Ardizzone, L., Kruse, J., Wirkert, S., Rahner, D., Pellegriani, E.W., Klessen, R.S., Maier-Hein, L., Rother, C., Köthe, U., 2018. Analyzing inverse problems with invertible neural networks. *arXiv preprint arXiv:1808.04730*.
- [3] Argun, A., Volpe, G., Bo, S., 2021. Classification, inference and segmentation of anomalous diffusion with recurrent neural networks. *Journal of Physics A: Mathematical and Theoretical* 54, 294003.
- [4] Arutkin, M., Walter, B., Wiese, K.J., 2020. Extreme events for fractional brownian motion with drift: Theory and numerical validation. *Physical Review E* 102. URL: <https://doi.org/10.1103/physreve.102.022102>, doi:doi:10.1103/physreve.102.022102.
- [5] Azevedo, T., Campbell, A., Romero-Garcia, R., Passamonti, L., Bethlehem, R.A., Liò, P., Toschi, N., 2020. A deep graph neural network architecture for modelling spatio-temporal dynamics in resting-state mri data. *bioRxiv*.
- [6] Bengio, Y., Courville, A., Vincent, P., 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence* 35, 1798–1828.
- [7] Bishop, C.M., . *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer. URL: <https://www.xarg.org/ref/a/0387310738/>.
- [8] Bo, S., Schmidt, F., Eichhorn, R., Volpe, G., 2019. Measurement of anomalous diffusion using recurrent neural networks doi:doi:10.1103/PhysRevE.100.010102, *arXiv:arXiv:1905.02038*.
- [9] Both, G.J., Kusters, R., 2021. Fully differentiable model discovery. *arXiv preprint arXiv:2106.04886*.
- [10] Bouchaud, J., 2003. *Theory of financial risk and derivative pricing : from statistical physics to risk management*. Cambridge University Press, Cambridge.
- [11] Bouchaud, J.P., Georges, A., 1990. Anomalous diffusion in disordered media: Statistical mechanisms, models and physical applications. *Physics Reports* 195, 127–293. URL: [https://doi.org/10.1016/0370-1573\(90\)90099-n](https://doi.org/10.1016/0370-1573(90)90099-n), doi:doi:10.1016/0370-1573(90)90099-n.
- [12] Burov, S., Jeon, J.H., Metzler, R., Barkai, E., 2011. Single particle tracking in systems showing anomalous diffusion: the role of weak ergodicity breaking. *Physical Chemistry Chemical Physics* 13, 1800. URL: <https://doi.org/10.1039/c0cp01879a>, doi:doi:10.1039/c0cp01879a.
- [13] Charles, R.Q., Su, H., Kaichun, M., Guibas, L.J., 2017. PointNet: Deep learning on point sets for 3d classification and segmentation, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE. URL: <https://doi.org/10.1109/cvpr.2017.16>, doi:doi:10.1109/cvpr.2017.16.
- [14] Cranmer, K., Brehmer, J., Louppe, G., 2020. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences* 117, 30055–30062. URL: <https://doi.org/10.1073/pnas.1912789117>, doi:doi:10.1073/pnas.1912789117.
- [15] Crispin, G., a. *Handbook of Stochastic Methods: for Physics, Chemistry and natural sciences*. 4th ed., Springer.
- [16] Crispin, G., b. *Handbook of Stochastic Methods: for Physics, Chemistry and natural sciences*. 4th ed., Springer.
- [17] Cutland, N.J., Kopp, P.E., Willinger, W., 1995. Stock price returns and the joseph effect: A fractional version of the black-scholes model, in: *Seminar on Stochastic Analysis, Random Fields and Applications*. Birkhäuser Basel, pp. 327–351. URL: https://doi.org/10.1007/978-3-0348-7026-9_23, doi:doi:10.1007/978-3-0348-7026-9_23.
- [18] Davies, R.B., Harte, D., 1987. Tests for hurst effect. *Biometrika* 74, 95–101.
- [19] Decreusefond, L., Üstünel, A.S., 1998. Fractional brownian motion: theory and applications. *ESAIM: Proceedings* 5, 75–86. URL: <https://doi.org/10.1051/proc:1998014>, doi:doi:10.1051/proc:1998014.
- [20] Defferrard, M., Bresson, X., Vandergheynst, P., 2016. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in neural information processing systems* 29.
- [21] Deng, W., Barkai, E., 2009. Ergodic properties of fractional brownian-langevin motion. *Physical Review E* 79. URL: <https://doi.org/10.1103/physreve.79.011112>, doi:doi:10.1103/physreve.79.011112.
- [22] Dinh, L., Sohl-Dickstein, J., Bengio, S., 2016. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*.
- [23] Du, S.S., Zhai, X., Poczos, B., Singh, A., 2018. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*.
- [24] Dubbeldam, J.L.A., Rostiashvili, V.G., Milchev, A., Vilgis, T.A., 2011. Fractional brownian motion approach to polymer translocation: The governing equation of motion. *Physical Review E* 83. URL: <https://doi.org/10.1103/physreve.83.011802>, doi:doi:10.1103/physreve.83.011802.
- [25] Ernst, D., Hellmann, M., Köhler, J., Weiss, M., 2012. Fractional brownian motion in crowded fluids. *Soft Matter* 8, 4886. URL: <https://doi.org/10.1039/c2sm25220a>, doi:doi:10.1039/c2sm25220a.
- [26] Fey, M., Lenssen, J.E., 2019. Fast graph representation learning with pytorch geometric. *arXiv:arXiv:1903.02428*.
- [27] Geneston, E., Tuladhar, R., Beig, M.T., Bologna, M., Grigolini, P., 2016. Ergodicity breaking and localization. *Physical Review E* 94. URL: <https://doi.org/10.1103/physreve.94.012136>, doi:doi:10.1103/physreve.94.012136.
- [28] Gerken, J.E., Aronsson, J., Carlsson, O., Linander, H., Ohlsson, F., Petersson, C., Persson, D., 2021. Geometric deep learning and equivariant neural networks. *arXiv preprint arXiv:2105.13926*.
- [29] Gherardi, M., Calabrese, L., Tamm, M., Lagomarsino, M.C., 2017. Model of chromosomal loci dynamics in bacteria as fractional diffusion with intermittent transport. *Physical Review E* 96. URL: <https://doi.org/10.1103/physreve.96.042402>, doi:doi:10.1103/physreve.96.042402.

- [30] Han, D., Korabel, N., Chen, R., Johnston, M., Gavrilova, A., Allan, V.J., Fedotov, S., Waigh, T.A., 2020. Deciphering anomalous heterogeneous intracellular transport with neural networks. *eLife* 9. URL: <https://doi.org/10.7554/eLife.52224>, doi:doi:10.7554/eLife.52224.
- [31] Javer, A., Kuwada, N.J., Long, Z., Benza, V.G., Dorfman, K.D., Wiggins, P.A., Cicuta, P., Lagomarsino, M.C., 2014. Persistent super-diffusive motion of escherichia coli chromosomal loci. *Nature Communications* 5. URL: <https://doi.org/10.1038/ncomms4854>, doi:doi:10.1038/ncomms4854.
- [32] Jeon, J.H., Metzler, R., 2012. Inequivalence of time and ensemble averages in ergodic systems: Exponential versus power-law relaxation in confinement. *Physical Review E* 85. URL: <https://doi.org/10.1103/physreve.85.021147>, doi:doi:10.1103/physreve.85.021147.
- [33] Jeon, J.H., Tejedor, V., Burov, S., Barkai, E., Selhuber-Unkel, C., Berg-Sørensen, K., Oddershede, L., Metzler, R., 2011. In Vivo Anomalous diffusion and weak ergodicity breaking of lipid granules. *Physical Review Letters* 106. URL: <https://doi.org/10.1103/physrevlett.106.048103>, doi:doi:10.1103/physrevlett.106.048103.
- [34] Kipf, T.N., Welling, M., 2016. Semi-supervised classification with graph convolutional networks. *arXiv:arXiv:1609.02907*.
- [35] Kobayev, I., Prince, S., Brubaker, M., 2020. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1. URL: <https://doi.org/10.1109/tpami.2020.2992934>, doi:doi:10.1109/tpami.2020.2992934.
- [36] Koo, P.K., Mochrie, S.G.J., 2016. Systems-level approach to uncovering diffusive states and their transitions from single-particle trajectories. *Physical Review E* 94. URL: <https://doi.org/10.1103/physreve.94.052412>, doi:doi:10.1103/physreve.94.052412.
- [37] Kosztolowicz, T., Dworecki, K., Mrówczyński, S., 2005. How to measure subdiffusion parameters. *Physical Review Letters* 94. URL: <https://doi.org/10.1103/physrevlett.94.170602>, doi:doi:10.1103/physrevlett.94.170602.
- [38] Krapf, D., Lukat, N., Marinari, E., Metzler, R., Oshanin, G., Selhuber-Unkel, C., Squarcini, A., Stadler, L., Weiss, M., Xu, X., 2019. Spectral content of a single non-brownian trajectory. *Physical Review X* 9, 011019.
- [39] Kraskov, A., Stögbauer, H., Grassberger, P., 2004. Estimating mutual information. *Physical review E* 69, 066138.
- [40] Krog, J., Jacobsen, L.H., Lund, F.W., Wustner, D., Lomholt, M.A., 2018. Bayesian model selection with fractional brownian motion. *Journal of Statistical Mechanics: Theory and Experiment* 2018. URL: <https://doi.org/10.1088/1742-5468/aadb0e>, doi:doi:10.1088/1742-5468/aadb0e.
- [41] Kukla, V., Kornatowski, J., Demuth, D., Girnus, I., Pfeifer, H., Rees, L.V.C., Schunk, S., Unger, K.K., Karger, J., 1996. NMR studies of single-file diffusion in unidimensional channel zeolites. *Science* 272, 702–704. URL: <https://doi.org/10.1126/science.272.5262.702>, doi:doi:10.1126/science.272.5262.702.
- [42] Kursawe, J., Schulz, J., Metzler, R., 2013. Transient aging in fractional brownian and langevin-equation motion. *Physical Review E* 88. URL: <https://doi.org/10.1103/physreve.88.062124>, doi:doi:10.1103/physreve.88.062124.
- [43] LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *nature* 521, 436–444.
- [44] Leland, W., Taqqu, M., Willinger, W., Wilson, D., 1994. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking* 2, 1–15. URL: <https://doi.org/10.1109/90.282603>, doi:doi:10.1109/90.282603.
- [45] Loch-Olszewska, H., Sikora, G., Janczura, J., Weron, A., 2016. Identifying ergodicity breaking for fractional anomalous diffusion: Criteria for minimal trajectory length. *Physical Review E* 94. URL: <https://doi.org/10.1103/physreve.94.052136>, doi:doi:10.1103/physreve.94.052136.
- [46] Lu, Y.J., Li, C.T., 2021. Agstn: Learning attention-adjusted graph spatio-temporal networks for short-term urban sensor value forecasting. *arXiv:arXiv:2101.12465*.
- [47] Lysy, M., Pillai, N.S., Hill, D.B., Forest, M.G., Mellnik, J.W.R., Vasquez, P.A., McKinley, S.A., 2016. Model comparison and assessment for single particle tracking in biological fluids. *Journal of the American Statistical Association* 111, 1413–1426. URL: <https://doi.org/10.1080/01621459.2016.1158716>, doi:doi:10.1080/01621459.2016.1158716.
- [48] Mandelbrot, B.B., Ness, J.W.V., 1968. Fractional brownian motions, fractional noises and applications. *SIAM Review* 10, 422–437. URL: <https://doi.org/10.1137/1010093>, doi:doi:10.1137/1010093.
- [49] McInnes, L., Healy, J., Melville, J., 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- [50] Meroz, Y., Sokolov, I.M., 2015. A toolbox for determining subdiffusive mechanisms. *Physics Reports* 573, 1–29. URL: <https://doi.org/10.1016/j.physrep.2015.01.002>, doi:doi:10.1016/j.physrep.2015.01.002.
- [51] Metzler, R., Jeon, J.H., Cherstvy, A.G., Barkai, E., 2014. Anomalous diffusion models and their properties: non-stationarity, non-ergodicity, and ageing at the centenary of single particle tracking. *Phys. Chem. Chem. Phys.* 16, 24128–24164. URL: <https://doi.org/10.1039/c4cp03465a>, doi:doi:10.1039/c4cp03465a.
- [52] Monahan, J.F., 2011. Numerical methods of statistics. Cambridge University Press.
- [53] Muñoz-Gil, G., Volpe, G., Garcia-March, M.A., Aghion, E., Argun, A., Hong, C.B., Bland, T., Bo, S., Conejero, J.A., Firdas, N., et al., 2021. Objective comparison of methods to decode anomalous diffusion. *Nature communications* 12, 1–16.
- [54] Muñoz-Gil, G., Volpe, G., Garcia-March, M.A., Metzler, R., Lewenstein, M., Manzo, C., 2020. The anomalous diffusion challenge: single trajectory characterisation as a competition. *Emerging Topics in Artificial Intelligence* 2020. URL: <http://dx.doi.org/10.1117/12.2567914>, doi:doi:10.1117/12.2567914.
- [55] Garibo-i Orts, Ò., Baeza-Bosca, A., Garcia-March, M.A., Conejero, J.A., 2021. Efficient recurrent neural network methods for anomalously diffusing single particle short and noisy trajectories. *Journal of Physics A: Mathematical and Theoretical* 54, 504002.
- [56] Palyulin, V.V., Ala-Nissila, T., Metzler, R., 2014. Polymer translocation: the first two decades and the recent diversification. *Soft Matter* 10, 9016–9037. URL: <https://doi.org/10.1039/c4sm01819b>, doi:doi:

- 10.1039/c4sm01819b.
- [57] Papamakarios, G., Sterratt, D., Murray, I., 2019. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows, in: Chaudhuri, K., Sugiyama, M. (Eds.), Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics, PMLR. pp. 837–848. URL: <http://proceedings.mlr.press/v89/papamakarios19a.html>.
- [58] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830.
- [59] Qi, C.R., Yi, L., Su, H., Guibas, L., 2017. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *ArXiv preprint arXiv:1706.02413*.
- [60] Radev, S.T., D’Alessandro, M., Mertens, U.K., Voss, A., Köthe, U., Bürkner, P.C., 2021. Amortized bayesian model comparison with evidential deep learning. *IEEE Transactions on Neural Networks and Learning Systems*.
- [61] Radev, S.T., Mertens, U.K., Voss, A., Ardizzone, L., Kothe, U., 2020. BayesFlow: Learning complex stochastic models with invertible neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 1–15 URL: <https://doi.org/10.1109/tnnls.2020.3042395>, doi:doi:10.1109/tnnls.2020.3042395.
- [62] Raissi, M., Perdikaris, P., Karniadakis, G.E., 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics* 378, 686–707.
- [63] Reverey, J.F., Jeon, J.H., Bao, H., Leippe, M., Metzler, R., Selhuber-Unkel, C., 2015. Superdiffusion dominates intracellular particle motion in the supercrowded cytoplasm of pathogenic *acanthamoeba castellanii*. *Scientific reports* 5, 1–14.
- [64] Rostek, S., Schöbel, R., 2013. A note on the use of fractional brownian motion for financial modeling. *Economic Modelling* 30, 30–35. URL: <https://doi.org/10.1016/j.econmod.2012.09.003>, doi:doi:10.1016/j.econmod.2012.09.003.
- [65] Sanders, L.P., Ambjörnsson, T., 2012. First passage times for a tracer particle in single file diffusion and fractional brownian motion. *The Journal of Chemical Physics* 136, 175103. URL: <https://doi.org/10.1063/1.4707349>, doi:doi:10.1063/1.4707349.
- [66] Thapa, S., Lomholt, M.A., Krog, J., Cherstvy, A.G., Metzler, R., 2018. 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, 29018–29037. URL: <https://doi.org/10.1039/c8cp04043e>, doi:doi:10.1039/c8cp04043e.
- [67] Thapa, S., Park, S., Kim, Y., Jeon, J.H., Metzler, R., Lomholt, M.A., 2022. Bayesian inference of scaled versus fractional brownian motion. *Journal of Physics A: Mathematical and Theoretical* 55, 194003.
- [68] Ulyanov, D., Vedaldi, A., Lempitsky, V., 2016. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*.
- [69] Verdier, H., Duval, M., Laurent, F., Casse, A., Vestergaard, C.L., Masson, J.B., 2021. Learning physical properties of anomalous random walks using graph neural networks. *Journal of Physics A: Mathematical and Theoretical* URL: <https://doi.org/10.1088/1751-8121/abfa45>, doi:doi:10.1088/1751-8121/abfa45.
- [70] Verdier, H., Laurent, F., Cassé, A., Vestergaard, C.L., Specht, C.G., Masson, J.B., 2022. A maximum mean discrepancy approach reveals subtle changes in α -synuclein dynamics. *bioRxiv*.
- [71] Walter, J.C., Ferrantini, A., Carlon, E., Vanderzande, C., 2012. Fractional brownian motion and the critical dynamics of zipping polymers. *Physical Review E* 85. URL: <https://doi.org/10.1103/physreve.85.031120>, doi:doi:10.1103/physreve.85.031120.
- [72] Wang, W., Cherstvy, A.G., Chechkin, A.V., Thapa, S., Seno, F., Liu, X., Metzler, R., 2020. Fractional brownian motion with random diffusivity: emerging residual non-ergodicity below the correlation time. *Journal of Physics A: Mathematical and Theoretical* 53, 474001. URL: <https://doi.org/10.1088/1751-8121/aba467>, doi:doi:10.1088/1751-8121/aba467.
- [73] Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M., 2018. Dynamic graph cnn for learning on point clouds. *arXiv:arXiv:1801.07829*.
- [74] Waterfall, J.J., Casey, F.P., Gutenkunst, R.N., Brown, K.S., Myers, C.R., Brouwer, P.W., Elser, V., Sethna, J.P., 2006. Sloppy-model universality class and the vandermonde matrix. *Physical review letters* 97, 150601.
- [75] Weber, S.C., Spakowitz, A.J., Theriot, J.A., 2010a. Bacterial chromosomal loci move subdiffusively through a viscoelastic cytoplasm. *Physical Review Letters* 104. URL: <https://doi.org/10.1103/physrevlett.104.238102>, doi:doi:10.1103/physrevlett.104.238102.
- [76] Weber, S.C., Spakowitz, A.J., Theriot, J.A., 2010b. Bacterial chromosomal loci move subdiffusively through a viscoelastic cytoplasm. *Physical review letters* 104, 238102.
- [77] Xu, K., Hu, W., Leskovec, J., Jegelka, S., 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*.
- [78] Ying, R., You, J., Morris, C., Ren, X., Hamilton, W.L., Leskovec, J., 2018. Hierarchical graph representation learning with differentiable pooling. *arXiv:arXiv:1806.08804*.
- [79] Yu, S., Wu, J., Meng, X., Chu, R., Li, X., Wu, G., 2021. Mesoscale simulation of bacterial chromosome and cytoplasmic nanoparticles in confinement. *Entropy* 23, 542. URL: <https://doi.org/10.3390/e23050542>, doi:doi:10.3390/e23050542.

SUPPLEMENTARY MATERIAL

b. GNN Architecture

1. amortised inference model architecture and training

a. Node and edge features

The features associated to nodes and edges of a trajectory's graph in the GRATIN summary network are defined using the following variables defined for each node $i \in \{1, \dots, N\}$: the sum $R_i^{(k)}$ of step sizes elevated to the power k , up to point i : $\sum_{j < i} \|\Delta r_j^k\|_2$. We denote as well s the standard deviation of step sizes. Then, the node features associated to a node i are:

1. the normalised time: i/N ,
2. the normalised distance to origin $\|r_i\|_2/s\sqrt{i}$,
3. the normalised maximal distance to origin up to point i : $\max_{k \leq i} \|r_k\|_2/s\sqrt{i}$,
4. normalised cumulative powered distances covered by the walker up to point i , with $k \in \{1, 2, 4\}$: $\frac{N}{i} \frac{R_i^{(k)}}{R_N^{(k)}}$.

The six features associated to an edge (i, j) with $i < j$ are:

1. the time difference: $j - i$,
2. the normalised distance between edge source and target: $\|r_j - r_i\|_2/s\sqrt{j - i}$,
3. the normalised dot product of jumps: $\Delta r_i^\top \Delta r_j/s^2$,
4. normalised differences of the cumulative powered distances covered by the walker from point i to j , with $k \in \{1, 2, 4\}$: $\frac{N}{j-i} \frac{R_j^{(k)} - R_i^{(k)}}{R_N^{(k)}}$.

The number of operations required to compute features scales linearly with the trajectory length.

The GRATIN-NF version of the summary network simply uses two features per node and edge. Its node features are :

1. the normalised time: i/N ,
2. the normalised distance to origin: $\|r_i\|_2/s\sqrt{i}$.

Its edge features are :

1. the normalised time difference : $(j - i)/N$,
2. the normalised distance : $\|r_i - r_j\|_2/s\sqrt{j - i}$.

The architecture of the GNN used in the summary network is similar to the encoder network proposed in [69], with the difference that we here additionally apply edge features. Node and edge features are first passed to perceptrons, which embed them in an 8-dimensional space. The network is then composed of three successive convolution layers, one relying simply on node features and the two others being conditioned by edge features. We used as convolutions the GIN layers introduced in [77]). They output $\mathbf{x}^{(1)}$, $\mathbf{x}^{(2)}$ and $\mathbf{x}^{(3)}$ vectors, each of 32 dimensions (equivalent to 32 convolution filters), which are concatenated to form $\mathbf{x}^{(f)}$. The rows of this $(N, 32 \times 3)$ matrix of nodes features are then aggregated using an attention mechanism during the pooling step, to keep just one row per graph, i.e., per trajectory. This vector is subsequently passed to a three-layer perceptron, the output of which is a 12-dimensional summary statistics vector. The dimension of the summary statistics vector is voluntarily higher than that of the parameters space so as to facilitate the neural network convergence, which is helped by over-parameterisation [23].

c. Differences with Verdier et al.

The network introduced here performs variational inference on a well-defined class of trajectories, while the one presented in [69] provides a point estimation of the anomalous exponent and the random walk type for a variety of random walk models, accounting for experimental noise. The main similarity between the two networks is the GNN structure of their encoders. Here is an extensive list of the differences between these two modules :

1. In the present work, features are not only associated to nodes but to edges as well.
2. We use a more expressive graph convolution layer, able to process edge features, the GIN layer introduced in [77].
3. We use batch normalisation layers which are more appropriate to graphs, the InstanceNorm layer introduced in [68].
4. The argument of the geometric series used to wire nodes is no more fixed by the length of the trajectory but instead depends on the index of the node of destination.

d. LSTM architecture

The summary network based on LSTM takes as input, for each trajectory, the vector of positions. It has the following sequence of layers:

- 2 1-dimensional convolutions, with 64 convolution kernels of size 8.
- 4 stacked bi-directional LSTM whose hidden units have 32 dimensions
- 1 MLP of size $(32 \times 2 \times 4, 12)$. The input dimension corresponds to the product of the number of dimension of hidden units, the number of stacked layers and the number of directions. Indeed, the MLP takes as input a concatenation of the eight last hidden states of the LSTMs. The output has the dimension of the latent space.

e. Invertible network

The invertible network is a succession of three affine coupling blocks. These blocks, introduced in [22], transform an input vector \mathbf{u} into \mathbf{v} in an invertible manner parameterised by the summary statistics vector \mathbf{h} . They do so by splitting \mathbf{u} into two halves \mathbf{u}_1 and \mathbf{u}_2 , used to compute the two halves \mathbf{v}_1 and \mathbf{v}_2 of \mathbf{v} by consecutively performing the two following operations :

$$\begin{aligned}\mathbf{v}_1 &= \mathbf{u}_1 \odot \exp(s_1(\mathbf{u}_2; \mathbf{h})) + t_1(\mathbf{u}_2; \mathbf{h}) \\ \mathbf{v}_2 &= \mathbf{u}_2 \odot \exp(s_2(\mathbf{v}_1; \mathbf{h})) + t_2(\mathbf{v}_1; \mathbf{h})\end{aligned}$$

where \odot denotes the element-wise multiplication (the Hadamard product) and where s_1, s_2, t_1 and t_2 are multi-layer perceptrons, which do not need to be invertible. In our case, they have five hidden layers and their activation function is an exponential linear unit. This procedure can be inverted to efficiently retrieve \mathbf{u} from \mathbf{v} .

f. Training the networks

Not all parameter directions in the parameter space are equally constrained by the data. Thus, we split the summary statistics vector in two halves (one per parameter to infer) and pre-train summary networks to infer each parameter individually. This is motivated by the expected misbehaviour of variational optimisation for an inference whose parameters are under significantly different constraints. Hence, the good performance at inferring an easily learnt parameter (such as K_α) does not prevent the network from converging towards a better optimum where it infers more challenging parameters as well. We do this by using the output of the encoder GNN as an input to a multi-layer perceptron, and optimising the so-obtained regressor to infer a given parameter in a regression setting. The multi-layer perceptron is then discarded, and the outputs of the parameter-specific GNNs are concatenated to form the summary statistics vector used in the coupled inference. Weights of the summary networks are then frozen and only the invertible part of the network is trained.

2. Exact posterior inference

To compute exact posteriors, likelihood values were computed on grids of points in parameter space. We picked a uniform prior on $(0.4, 1.6)$ for α and a log-uniform one for τ_c and K_α , which spanned 8 orders of magnitude. There was thus no coupling between parameters in the priors. The parameters used to generate trajectories during training were sampled from these same priors.

3. Cramér-Rao bound

Formally, we consider any estimator of the parameters $\boldsymbol{\theta}$ to be a (possibly implicit) function of the recorded trajectory, \mathbf{R} , i.e., $\hat{\boldsymbol{\theta}} = \mathbf{T}(\mathbf{R})$. We denote by $\boldsymbol{\psi}(\boldsymbol{\theta}) = E[\mathbf{T}(\mathbf{R})]$ its expectation, and by $\Gamma(\boldsymbol{\theta}) = E[(\mathbf{T}(\mathbf{R}) - \boldsymbol{\psi}(\boldsymbol{\theta}))(\mathbf{T}(\mathbf{R}) - \boldsymbol{\psi}(\boldsymbol{\theta}))^\top]$ its covariance matrix. Finally, $\mathbf{I}(\boldsymbol{\theta})$ is the Fischer information matrix, whose elements are given by $\mathbf{I}_{n,m}(\boldsymbol{\theta}) = E\left[\frac{\partial}{\partial \theta_n} \log p(\mathbf{R}|\boldsymbol{\theta}) \frac{\partial}{\partial \theta_m} \log p(\mathbf{R}|\boldsymbol{\theta})\right]$.

The Cramér-Rao bound states that, for any estimator \mathbf{T} ,

$$\Gamma(\boldsymbol{\theta}) \geq \nabla \boldsymbol{\psi}(\boldsymbol{\theta}) [\mathbf{I}(\boldsymbol{\theta})]^{-1} [\nabla \boldsymbol{\psi}(\boldsymbol{\theta})]^\top, \quad (6)$$

where $\nabla \boldsymbol{\psi}$ is the Jacobian of $\boldsymbol{\psi}$. In particular, this matrix inequality implies the following lower bound on the variance of any unbiased estimator of a single parameter:

$$\text{Var}_{\boldsymbol{\theta}}(T_n(\mathbf{R})) \geq [\mathbf{I}(\boldsymbol{\theta})^{-1}]_{n,n} \quad (7)$$

. This inequality is loose and the bound exaggeratedly low if parameters are not approximately equally difficult to estimate. Hence, when computing the lower bound of the variance $b_N(\alpha)$, we treated the problem of section III as univariate, i.e. we discarded the terms related to the inference of K_α in the Fischer information matrix.

4. Supplementary figures




FIG. S1. **Latent space representations of individual trajectories.** 2D visualisation of summary vectors (one point per trajectory), obtained by UMAP and colored according to A: their anomalous diffusion exponent α , and B: their correlation time τ_c . Trajectories are of length $N = 1,000$.




FIG. S2. **Robustness to noise.** MSE on α estimated with amortised inference compared to the Cramér-Rao bound as a function of positioning noise σ . Trajectories are of length $N = 200$ and generalised diffusivity 1. Positions are independently corrupted with Gaussian noise of variance σ^2 . The diffusion coefficient was set to 1, α was sampled from a uniform distribution between 0.4 and 1.6, and the performance was obtained by averaging over 10 000 trajectories for each level of noise




FIG. S3. **Temporal correlations of fBm with finite decorrelation time.** Autocovariance of increments of the fBm trajectory, with finite correlation time τ_c , in the sub-diffusive and super-diffusive case. Red curves correspond to $\alpha = 0.6$, and $\tau_c = 5$ (dashed line), 15 (dotted line), ∞ (plain line). Green curves correspond to $\alpha = 1.4$, and $\tau_c = 5, 25, \infty$.