

Planted directed polymer: Inferring a random walk from noisy images

Sun Woo P. Kim ^{*}

Department of Physics, King's College London, Strand, London WC2R 2LS, United Kingdom

Austen Lamacraft 

TCM Group, Cavendish Laboratory, University of Cambridge, Cambridge CB3 0HE, United Kingdom



(Received 29 June 2024; revised 4 December 2024; accepted 7 February 2025; published 20 February 2025)

We introduce and study the planted directed polymer, in which the path of a random walker is inferred from noisy "images" accumulated at each time step. Formulated as a nonlinear problem of Bayesian inference for a hidden Markov model, this problem is a generalization of the directed polymer problem of statistical physics, coinciding with it in the limit of zero signal to noise. For a one-dimensional walker we present numerical investigations and analytical arguments that no phase transition is present. When formulated on a Cayley tree, methods developed for the directed polymer are used to show that there is a transition with decreasing signal to noise where effective inference becomes impossible, meaning that the average fractional overlap between the inferred and true paths falls from one to zero.

DOI: [10.1103/PhysRevE.111.024135](https://doi.org/10.1103/PhysRevE.111.024135)

I. INTRODUCTION

Recent years have seen a great deal of research activity at the interface of statistical physics and Bayesian inference [1]. Broadly speaking, the connection is as follows. Suppose the system of interest is described by some random variables $\mathbf{x}$ following a prior probability distribution $p(\mathbf{x})$. We learn about $\mathbf{x}$ by measuring some other random variables $\mathbf{y}$, which are described by some conditional distribution or "measurement model" $p(\mathbf{y}|\mathbf{x})$ [2]. From the values of $\mathbf{y}$ we compute the posterior distribution $p(\mathbf{x}|\mathbf{y})$ of $\mathbf{x}$ using Bayes' rule

$$p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})}, \quad (1)$$

where the distribution of the measurement outcomes $p(\mathbf{y}) = \sum_{\mathbf{x}} p(\mathbf{y}|\mathbf{x})p(\mathbf{x})$ can be regarded as a normalizing factor or partition function in the language of statistical physics. This setting assumes a perfect knowledge of the prior distribution and measurement model, an idealization which could be approached in certain (highly engineered) settings, e.g., transmission over a noisy channel. Nevertheless, the computation of $p(\mathbf{x}|\mathbf{y})$ will in general be intractable when the number of variables $\mathbf{x}$ is large, due to the difficulty of calculating the denominator in Eq. (1). The intractability of the partition function is a general feature of statistical mechanical systems, and this represents the first point of contact between statistical physics and inference. Furthermore, there are "natural" examples of inference problems in the "thermodynamic limit" of

many variables where, for certain cases, the phenomenology of phase transitions in the thermodynamic limit of random systems is applicable, with the values of the measurements $\mathbf{y}$ corresponding to the randomness. In the inference setting, such transitions represent a change in the structure of the posterior as a parameter representing the strength of the "signal" of $\mathbf{x}$ in the measurements $\mathbf{y}$ is varied. In cases where there exists a phase transition, as the signal weakens, the probability of successful inference goes to zero at the phase transition.

In this work, we will be concerned with developing the parallels between statistical mechanics and inference within a particular class of problems that have the structure of a hidden Markov model (HMM; Fig. 1) [3,4]. In an HMM, the state of the system x_t follows a Markov process,

$$p(x_{1:t}) = \left(\prod_{\tau=2}^t p(x_{\tau}|x_{\tau-1}) \right) p(x_1), \quad (2)$$

where we have introduced the notation $x_{1:t}$ to denote a sequence of variables $x_1, \dots, x_t$, and $p(x_{\tau}|x_{\tau-1})$ are the transition probabilities (or kernel) of the process. Measurements y_t are conditional on the value x_t at the same time step through some measurement model $p(y_t|x_t)$.

Inference in HMMs may involve one of several related tasks. Filtering refers to the problem of obtaining the posterior $p(x_t|y_{1:t})$ of the present value conditioned on the history of measurements, while smoothing involves conditioning additionally over future measurements. The most famous example of filtering in a HMM is the Kalman filter [5], in which both $p(x_{\tau}|x_{\tau-1})$ and $p(x_t|y_t)$ are Gaussian with linear dependence on the conditioning variable. Because the product of Gaussians is also Gaussian, the filtering distribution $p(x_t|y_{1:t})$ —obtained from Bayes' rule Eq. (1)—is a Gaussian with an explicit form.

^{*}Contact author: swk34@cantab.ac.uk

Here we are concerned with a fundamentally nonlinear observation model, consisting of "images" containing "pixels," one for each possible location of the system x . This is a very natural model of tracking an object from a time series of images taken against a noisy background, e.g., tracking molecules from fluorescence imaging [6]. From now on, we denote observations as $\phi_{x,t}$, indexed by both x and t . They take the form

$$\phi_{x,t} = \psi_{x,t} + \epsilon \delta_{x,x_t}, \quad (3)$$

where $\psi_{x,t}$ is Gaussian white noise, ϵ is a parameter that controls the relative strength of the signal (i.e., the signal-to-noise ratio), and δ_{x,x_t} is the Kronecker delta between pixel location x and current walker location x_t . In Eq. (3) the shift of the mean of $\phi_{x,t}$ by ϵ can reveal through which pixel the system has passed.

Problems of this type have appeared before in the engineering literature applied to tracking of paths through a noisy or cluttered environment [7,8]. Our aim is to introduce and analyze a particularly simple version of the problem that we call the *planted directed polymer problem*, because of its connection with the statistical mechanics of a directed polymer in a random medium [9].

A second and quite different motivation for our work is the recently discovered *measurement induced phase transition* in monitored quantum systems, where the entanglement dynamics and purity of extensive quantum systems was found to undergo a transition in the long-time limit as the measurement rate increases (see Refs. [10,11] and the review Ref. [12]). While there is as of yet no formal connection between the phenomena explored here and those in the quantum setting, the resemblance between the two situations is suggestive.

The outline of the rest of this paper is as follows. In the next section we introduce the planted directed polymer problem and describe the relation to the directed polymer. In Sec. III we study the one-dimensional case numerically to establish some general features of the phenomenology of the problem. In Sec. IV we present analytical results for the case of paths on a tree, building on the approach of Ref. [13] to the random polymer problem. Finally, we conclude in Sec. V and offer a perspective on similar problems in this class.

II. STATEMENT OF THE PLANTED DIRECTED POLYMER PROBLEM

In this section we introduce the planted directed polymer problem. We describe the simplest case where the dynamics of our system is a random walk on the integers $\mathbb{Z}$: the generalization of the model to more complicated situations, including higher dimensions, should be obvious. The position of the walker at (integer) time t is denoted by $x_t^* \in \mathbb{Z}$: throughout this work we follow Ref. [1] in using x_t^* to denote the "ground truth" or "planted" configuration, with x_t denoting the inferred configuration. The dynamics are specified by some transition kernel $p_T(x_{t+1}^* | x_t^*)$, with initial distribution $p_T(x_1^*)$ (the meaning of the subscript "T" will be explained shortly). As our main example we will take a symmetric random walk [see Fig. 2(a)]

$$p_T(x_{t+1}^* | x_t^*) = \frac{1}{2} \delta_{x_{t+1}^*, x_t^*} + \frac{1}{4} \delta_{x_{t+1}^*, x_t^* \pm 1}. \quad (4)$$




FIG. 1. Inference in a hidden Markov model.

Our goal is to infer the walker's trajectory from measurements that form noisy "images" specified by a "pixel" value $\phi_{x,t}$ at each position and time. These values are given by

$$\phi_{x,t} = \psi_{x,t} + \epsilon_T \delta_{x,x_t^*}, \quad (5)$$

where the pixel noise is a set of independent and identically distributed (i.i.d.) Gaussian distributions with standard deviation σ_T , and ϵ_T is the signal strength. An example is shown in Fig. 2(b). We will sometimes use the standard probabilists' notation $\mathcal{N}(\mu, \sigma^2)$ for a normal distribution of mean μ and standard deviation σ , so that $\phi_{x,t} \sim p(\cdot | x_t^*) = \mathcal{N}(\epsilon_T \delta_{x,x_t^*}, \sigma_T^2)$. We use the bold symbol $\boldsymbol{\phi}_t := (\phi_{x,t} | x \in \mathbb{Z})$ for the entire image at some time t . Following [1], we describe this process as the "teacher," and indicate its parameters using the subscript "T."

Given $\phi_{x,t}$, the "student" can perform Bayesian inference to obtain the posterior over the walker's trajectory $x_{1:t}$ given the past history of noisy images, $\boldsymbol{\phi}_{1:t} := (\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_t)$. We introduce the notation $X := x_{1:t}$, $X^* := x_{1:t}^*$ for the entire in-




FIG. 2. An example of Bayesian inference on a random walker from images with signal strength $\epsilon_T = 1.5$ and all other parameters $\sigma_T = \sigma_S = \epsilon_S = 1$. (a) is a realization of a random walker undergoing Brownian motion. (b) shows the sets of "noisy images" taken at each time step. (c) is the posterior distribution inferred from the noisy images.

ferred and planted configurations, respectively, and $\Phi := \phi_{1:t}$ for the sequence of images. The joint distribution is then

$$p(X, \Phi, X^*) = p_S(X|\Phi)p_T(\Phi|X^*)p_T(X^*). \quad (6)$$

$p_T(\Phi|X^*)$ is the probability of obtaining the sequence of images conditioned on the path X^* , given by the above Gaussian distribution, and the student's posterior over the path $X = x_{1:t}$ is given by Bayes' rule

$$p_S(X|\Phi) = \frac{p_S(\Phi|X)p_S(X)}{\sum_{X'} p_S(\Phi|X')p_S(X')}. \quad (7)$$

We use the subscript "S" to denote the student's distribution, which will allow us to consider a non-Bayes-optimal setting in which the parameters of the student's model do not coincide with those of the teacher. Later, we will also consider the case where the student also does not have access to the functional form of the teacher model. Equation (6) defines the *planted ensemble* (in the terminology of Ref. [1]) for our problem.

Using the rules of conditional probability on the student's prior we can also write the joint probability as

$$p(X, \Phi, X^*) = \frac{p_S(\Phi|X)p_S(X)p_T(\Phi|X^*)p_T(X^*)}{p_S(\Phi)}, \quad (8)$$

where $p_S(\Phi) = \sum_X p_S(\Phi|X)p_S(X)$. We see that in the Bayes optimal case, which is when the student and teacher distributions coincide, the joint distribution is symmetric with respect to $X \leftrightarrow X^*$. This means that inferred paths X must have the same distribution as that of the true paths X^* . In the statistical physics of inference literature, this is also known as the Nishimori condition.

Equation (6) represents a complete description of the problem, but dealing with the distribution over the entire trajectory is not normally practical. Instead, it is more usual to focus on the posterior distribution of the present value given the measurement history, i.e., $p_S(x_t|\phi_{1:t})$ (the filtering problem). Below we present a tractable algorithm to calculate the posterior for the present value.

At $t = 1$ Bayes' rule gives

$$p_S(x_1|\phi_1) = \frac{p_S(\phi_1|x_1)p_S(x_1)}{\sum_{x'_1} p_S(\phi_1|x'_1)p_S(x'_1)}. \quad (9)$$

$p_S(x_1)$ is the prior distribution on the initial position. $p_S(\phi_1|x_1)$ is the Gaussian distribution that follows from Eq. (5). At subsequent times, the student uses the posterior of the previous time steps to create the prior for the current time step:

$$p_S(x_t|\phi_{1:t}) = \frac{p_S(\phi_t|x_t)p_S(x_t|\phi_{1:t-1})}{\sum_{x'_t} p_S(\phi_t|x'_t)p_S(x'_t|\phi_{1:t-1})}, \quad (10)$$

where $p_S(x_t|\phi_{1:t-1})$ is given by the forward algorithm [14]

$$p_S(x_t|\phi_{1:t-1}) = \sum_{x_{t-1}} p_S(x_t|x_{t-1})p_S(x_{t-1}|\phi_{1:t-1}), \quad (11)$$

and $p_S(x_t|x_{t-1})$ is the transition kernel assumed by the student. An example of the posterior obtained in this way is depicted in Fig. 2(c).

Alternatively, it is often convenient to deal with an *unnormalized* posterior $q_S(x_t|\phi_{1:t})$ that obeys the linear equation

$$q_S(x_t|\phi_{1:t}) = q_S(\phi_t|x_t) \sum_{x_{t-1}} q_S(x_t|x_{t-1})q_S(x_{t-1}|\phi_{1:t-1}), \quad (12)$$

involving the unnormalized analogs of the transition and measurement probabilities. The right-hand sides of Eqs. (7) and (10) are both valid in terms of these unnormalized quantities, e.g., in Eq. (10) the likelihood $p_S(\phi_t|x_t)$ can be replaced with an unnormalized $q_S(\phi_t|x_t)$.

We are interested in distinguishing those regions in parameter space where the inferred trajectory x_t closely tracks the planted configuration x_t^* from those where it does not. We therefore consider the joint distribution

$$p(x_t, \phi_{1:t}, x_{1:t}^*) = p_S(x_t|\phi_{1:t})p_T(\phi_{1:t}|x_{1:t}^*)p_T(x_{1:t}^*). \quad (13)$$

We have already discussed the first factor in Eq. (10). The second factor corresponds to our Gaussian measurement model, Eq. (5). Finally, the third factor is the probability of a trajectory from the teacher Markov process:

$$p_T(x_{1:t}^*) = \left(\prod_{\tau=2}^t p_T(x_\tau^*|x_{\tau-1}^*) \right) p_T(x_1^*). \quad (14)$$

We can quantify the success of inference via the root-mean-squared error (RMSE) between the inferred position and true position at time t :

$$\text{RMSE}(t) := \sqrt{\mathbb{E}_{x_t, x_t^*} [(x_t - x_t^*)^2]}, \quad (15)$$

which generalizes in an obvious way to the case of d dimensions. Alternatively, we may consider the path-averaged quantity

$$\overline{\text{RMSE}}(t) = \sqrt{\frac{1}{t} \mathbb{E}_{X, X^*} \left[\sum_{\tau=1}^t (x_\tau - x_\tau^*)^2 \right]}. \quad (16)$$

If $\text{RMSE}(t)$ saturates at long times, $\text{RMSE}(\infty)$ and $\overline{\text{RMSE}}(\infty)$ will coincide.

In discrete settings we can also consider the average *overlap*: the fraction of the time that the inferred path coincides with the true path:

$$\bar{Y} := \frac{1}{t} \mathbb{E}_{X, X^*} \left[\sum_{\tau=1}^t \delta_{x_\tau, x_\tau^*} \right]. \quad (17)$$

$\bar{Y} = 1$ corresponds to perfect inference and $\bar{Y} \rightarrow 0$ at long times corresponds to failure of inference.

Relation to the directed polymer problem

Our Gaussian measurement model, Eq. (5), gives the likelihood of an image ϕ_t given current position x_t (for either student "S" or teacher "T" variables):

$$p(\phi_t|x_t) = \prod_{x'_t} \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{(\phi_{x'_t,t} - \epsilon \delta_{x_t, x'_t})^2}{2\sigma^2} \right] \quad (18)$$

$$= \exp \left[(\epsilon \phi_{x_t,t} - \epsilon^2/2)/\sigma^2 \right] \pi(\phi_t), \quad (19)$$

where $\pi(\phi_t)$ is the Gaussian measure with standard deviation σ over all spatial points. Since the image at a given time is

only conditionally dependent on the position of the walker at that time, the likelihood for all images factorizes as $p(\Phi|X) = \prod_{\tau=1}^t p(\phi_\tau|x_\tau)$. As x_t only enters in the Boltzmann-like factor, and recalling that the posterior Eq. (7) for the path $X = x_{1:t}$ can be written as a function of unnormalized prior $q(x_{1:t})$ and unnormalized likelihood,

$$p(X|\Phi) = \frac{\exp\left[\frac{\epsilon}{\sigma^2} \sum_{\tau=1}^t \phi_{x_\tau, \tau}\right] q(x_{1:t})}{Z(\Phi)}. \quad (20)$$

The normalizing factor, or partition function, is

$$Z(\Phi) := \sum_X \exp\left[\frac{\epsilon}{\sigma^2} \sum_{\tau=1}^t \phi_{x_\tau, \tau}\right] q(x_{1:t}). \quad (21)$$

The partition function can be rewritten in a form that is more reminiscent of the directed polymer in a random medium,

$$Z(\Phi) = \sum_X \exp\left[\sum_{\tau=1}^t h(x_{\tau-1:\tau}, \phi_\tau)\right] q(x_1), \quad (22)$$

with $h(x_{\tau-1:\tau}, \phi_\tau) := \ln q(x_\tau|x_{\tau-1}) + \frac{\epsilon}{\sigma^2} \phi_{x_\tau, \tau}$, corresponding to the "energy" of the polymer. The first term then corresponds to elasticity and the second to a potential. Here we set $q(x_1|x_0) = 1$ for convenience of notation.

Note that we can marginalize over the true paths X^* in the joint distributions, Eqs. (13) and (6), and instead consider $p(X, \Phi) = p_S(X|\Phi)p_T(\Phi)$, where the "planted" disorder is distributed as

$$p_T(\Phi) = \sum_{X^*} p_T(\Phi|X^*) p_T(X^*) \propto \pi_T(\Phi) Z_T(\Phi), \quad (23)$$

where $\pi_T(\Phi) = \prod_t \pi_T(\phi_t)$. Equation (23) shows how the Gaussian measure $\pi_T(\Phi)$ is modified by the presence of the ensemble of planted trajectories. When $\epsilon_T = 0$, $p_T(\Phi) = \pi_T(\Phi)$. In this case the joint probability $p(X, \Phi) = p_S(X|\Phi)\pi_T(\Phi)$ corresponds to that of the directed polymer in a random medium [9]. In terms of our inference problem, $\epsilon_T = 0$ means that the teacher provides no signal to the student, though the student believes a signal is present and uses $\epsilon_S \neq 0$ to construct the posterior $p_S(X|\Phi)$. As a result the path trajectory $x_{1:t}$ may be "pinned" in a random configuration by the "potential" $\phi_{x,t}$. When $\epsilon_T \neq 0$ the teacher provides a signal to the student (cf. Fig. 2) which alters the character of the posterior, as we will show in the next section.

Another way to view the problem is to separate Φ into random noise $\Psi := \psi_{1:t}$ and signal ϵ_T . We will see that this form is more helpful for theoretical analysis. In this case the posterior becomes conditional on both Ψ and X^* , as

$$p_S(X|\Psi, X^*) = \frac{\exp\left[\sum_{\tau=1}^t h_P(x_{\tau-1:\tau}, \psi_\tau, x_\tau^*)\right] q_S(x_1)}{Z_P(\Psi, X^*)}, \quad (24)$$

where $h_P(x_{\tau-1:\tau}, \psi_\tau, x_\tau^*) := \ln q_S(x_\tau|x_{\tau-1}) + \frac{\epsilon_S}{\sigma_S^2} \psi_{x_\tau, \tau} + \frac{\epsilon_S \epsilon_T}{\sigma_S^2} \delta_{x_\tau, x_\tau^*}$ and the *planted* partition function is

$$Z_P(\Psi, X^*) := \sum_X \exp\left[\sum_{\tau=1}^t h_P(x_{\tau-1:\tau}, \psi_\tau, x_\tau^*)\right] q_S(x_1). \quad (25)$$




FIG. 3. Plot of root-mean-squared error between true and inferred paths over time t for $\epsilon_S = \sigma_T = \sigma_S = 1$. As teacher signal strength ϵ_T increases, the root-mean-squared error plateaus at smaller values.

Then $p(\Psi) = \pi_T(\Psi)$ is just the i.i.d. Gaussian measure with teacher's parameters.

III. ONE-DIMENSIONAL CASE

In this section we present the results of numerical investigations of the properties of the posterior $p_S(X|\Phi)$ for a one-dimensional walker, and then further develop the above theoretical picture.

A. Numerical investigation

We solve Eq. (10) numerically using ϕ_t generated using $p_T(\phi_t|x_t^*)$ with x_t^* sampled from a symmetric random walk. We set $\epsilon_S = \sigma_T = \sigma_S = 1$ unless otherwise specified. Figure 3




FIG. 4. Scaling of estimated long-time $\text{RMSE}(\infty)$ with teacher signal strength ϵ_T for teacher noise strengths $\sigma_T = 0$ and $\sigma_T = 1$. Both cases support power law with ϵ_T , indicating a finite RMSE with any signal. The numerical results for the $\sigma_T = 0$ case confirms the $\propto \epsilon_T^{-1}$ theoretical prediction, Eq. (41). The $\sigma_T = 1$ case is also consistent with $\propto \epsilon_T^{-2}$ from our quantitative argument, Eq. (39).




FIG. 5. Conjectured phase diagram for wandering exponents α in the ϵ_S - ϵ_T plane with fixed $\sigma_S = \sigma_T = \sigma$. The blue and red texts are analytical results. Green is the conjectured behavior throughout the plane excluding $\epsilon_T = 0$. Numerics appear to support that $\alpha = 1/2$ for any finite $\epsilon_T > 0$.

shows the RMSE, Eq. (15), between the true and inferred positions—RMSE(t) with maximum time $t_{\max} = 1000$, system size $L = 1000$ with periodic boundary conditions—as ϵ_T is varied. We see that for ϵ_T close to 1, RMSE(t) appears to plateau within t_{plateau} at some RMSE(∞). In Fig. 4, for the values of ϵ_T where we can see a clear plateau, we study the scaling of estimated long-time RMSE(∞) with ϵ_T . For $\sigma_T = 1$ we see good evidence supporting that $\text{RMSE}(\infty) \sim \epsilon_T^{-2.06(1)}$, which, if the scaling holds for all values of $\epsilon_T > 0$, suggests that RMSE(t) is bounded for any $\epsilon_T > 0$. This suggests that

$\alpha = 1/2$ for any finite $\epsilon_T > 0$. For the $\sigma_T = 0$ case, the numerics support the $\text{RMSE}(\infty) \sim \epsilon_T^{-1}$ scaling of Eq. (41). This suggests that there is a change in the scaling of RMSE(∞) with and without the presence of disorder. We summarize our discussion on the wandering exponents in a conjectured phase diagram in Fig. 5.

In the directed polymer problem $\epsilon_T = 0$, there is no transition at finite noise strength σ_T , and the wandering exponent is $\alpha = 2/3$ for any $\sigma_T > 0$. Our numerics seem to suggest that the wandering exponent switches immediately to $\alpha = 1/2$ for any finite $\epsilon_T > 0$. In the language of directed polymers, $\epsilon_T > 0$ corresponds to the strength of an attractive delta potential. Our numerics suggest a hierarchy of terms in the polymer's energy Eq. (28): The disordered media term always overcomes the polymer "elasticity" term, while the delta-potential term always beats the disordered media term. In the language of Bayesian inference, inference is always possible with any true signal strength.

B. Directed polymer mapping

Here, to help with theoretical analysis, we take the continuum limit. Let us work with the separated noise and signal picture as Eq. (24). We define $\tilde{x} = ax$, $\tilde{\tau} = a\tau$, with final time $\tilde{t} = at$, then take the $a \rightarrow 0$ limit with $\tilde{t}$ fixed. Let us assume that the transition kernel $q_S(x_t|x_{t-1})$ only depends on the distance between the two points, and that it is maximized when the distance is zero. We take the scaling form $\ln q_S(x_t|x_{t-1}) = f_S(\sqrt{a}[x_t - x_{t-1}])$ and then expand around 0. We write $\psi_{x,t} = \sigma_T \eta_{x,t}$, where $\eta_{x,t} \sim \mathcal{N}(0, 1)$. Hence we can write the unnormalized posterior as

$$q_S(x_t|\psi_{1:t}, x_{1:t}^*) = \sum_{x_{t-1} \dots x_1} \exp \left(\sum_{\tau=1}^t \left[-\frac{a}{4\nu_S} [x_\tau - x_{\tau-1}]^2 + \frac{\epsilon_S \sigma_T}{\sigma_S^2} \eta_{x_\tau, \tau} + \frac{\epsilon_S \epsilon_T}{\sigma_S^2} \delta_{x_\tau, x_\tau^*} \right] + O(a^{3/2}) \right) q_S(x_1), \quad (26)$$

where $\nu_S = \frac{1}{2[f_S''(0)]}$ is the student's width of the kernel, and $x_0 = x_1$ for notational convenience. In the $a \rightarrow 0$ limit we have $\sum_\tau a \rightarrow \int d\tau$, $\frac{1}{a} \delta_{x_t, x_t^*} \rightarrow \delta[\tilde{x}(\tilde{t}) - \tilde{x}^*(\tilde{t})]$, and $\frac{ax_\tau - ax_{\tau-1}}{a} \rightarrow \frac{d\tilde{x}}{d\tilde{t}}$. We define $\tilde{\eta}_{x,t} := \frac{1}{a} \eta_{x,t}$ to retrieve white noise at continuum, $\mathbb{E}_{\tilde{\eta}} \tilde{\eta}_{x,t} \tilde{\eta}_{x',t'} = \frac{1}{a} \delta_{x,x'} \frac{1}{a} \delta_{t,t'} \rightarrow \delta(\tilde{x} - \tilde{x}') \delta(\tilde{t} - \tilde{t}')$. Taking the limit and then relabeling the continuum variables by removing the tildes, we have

$$q_S(x_t|\eta, x^*) \sim \int_{p_S[x(0)]}^{x(t)=x_t} D[x] e^{-E[x, \eta, x^*]}, \quad (27)$$

where now η and x^* refer to functions or continuous paths, and the energy of the polymer is

$$E[x, \eta, x^*] = \int_0^t d\tau \left[\frac{1}{4\nu_S} \dot{x}^2 - \frac{\epsilon_S \sigma_T}{\sigma_S^2} \eta(x(\tau), \tau) - \frac{\epsilon_S \epsilon_T}{\sigma_S^2} \delta[x(\tau) - x^*(\tau)] \right]. \quad (28)$$

This is the directed polymer with an attractive moving delta potential located at the path of the random walker $x^*(t)$. $p_S[x(0)]$ indicates the initial prior distribution. For example, for the Kronecker-delta initial state $p_S(x_1) = \delta_{x_1, 0}$ we can replace $p_S[x(0)] \rightarrow x(0) = 0$.

Taking the same limit for the partition function, the normalized posterior can then be written as

$$p_S(x_t|\eta, x^*) = \frac{q_S(x_t|\eta, x^*)}{Z_P(t)}, \quad (29)$$

where $Z_P(t) = \int_{p_S[x(0)]}^{x(t)=x_t^*} D[x] e^{-E[x, \eta, x^*]}$. In the continuum limit, we have $p_T(x_{1:t}^*) \sim e^{-E_0[x^*]} p_T[x^*(0)]$ where $E_0 = \int_0^t d\tau \frac{1}{4\nu_T} (\dot{x}^*)^2$. Marginalizing over $\eta_{1:t}$ and $x_{1:t-1}^*$ and then taking the continuum limit altogether, the joint distribution for inferred position and true position at time t is given by

$$p(x_t, x_t^*) = \int D\pi[\eta] \int_{p_T[x^*(0)]}^{x^*(t)=x_t^*} D[x^*] e^{-E_0[x^*]} \frac{q_S(x_t|\eta, x^*)}{Z_P(t)}. \quad (30)$$




FIG. 6. An example of the free-energy profile $F = -\ln p_S(x_t|\phi_{1:t})$ at $t = 1000$ for (a) delta peak prior and (b) flat prior. Vertical dashed pink line indicates true position x^* and dashed gray line indicates Gaussian modulation due to the delta peak prior.

Just as the directed polymer can be mapped to the Kardar-Parisi-Zhang (KPZ) equation and a diffusion process with random sources and sinks, from the path integral picture, we can see that the unnormalized posterior, with abuse of notation $q_S(x, t) = q_S(x_t|\eta, x^*)$, should evolve as a random diffusion process with a delta potential,

$$\frac{\partial q_S}{\partial t} = \nu_S \partial_x^2 q_S + \left(\frac{\epsilon_S \sigma_T}{\sigma_S^2} \eta(x, t) + \frac{\epsilon_S \epsilon_T}{\sigma_S^2} \delta[x - x^*(t)] \right) q_S, \quad (31)$$

and that the "free energy" $F := -\ln q_S(x, t)$ evolves as a KPZ equation with a delta potential,

$$\begin{aligned} \frac{\partial F}{\partial t} &= \nu_S \partial_x^2 F - \nu_S (\partial_x F)^2 - \frac{\epsilon_S \sigma_T}{\sigma_S^2} \eta(x, t) \\ &\quad - \frac{\epsilon_S \epsilon_T}{\sigma_S^2} \delta[x - x^*(t)]. \end{aligned} \quad (32)$$

In Fig. 6, we plot a typical free-energy profile for a given random walk x^* starting from both flat and delta initial priors $p_S(x_1) \sim 1$ and $\delta_{x_1, 0}$, respectively. We see that, superimposed on the rough landscape corresponding to KPZ behavior, there is a linear "wedge" in the free energy. In the noiseless case, $\sigma_T = 0$, Eq. (31) is the imaginary-time Schrödinger equation for a time-varying Hamiltonian with a delta potential located at x^* , $\hat{H}(t) = -\nu_S \partial_x^2 - \frac{\epsilon_S \epsilon_T}{\sigma_S^2} \delta[x - x^*(t)]$. The instantaneous ground state of $\hat{H}(t)$ is the single bound state given by $q_0(x) \sim e^{-\kappa|x-x^*|}$, with $\kappa = \frac{\epsilon_S \epsilon_T}{2\nu_S \sigma_S^2}$ and energy $E_0 = -\nu_S \kappa^2$. The imaginary-time evolution projects onto the ground state, and if the potential moves slowly in time, we expect the $q_S(x, t)$ to follow the moves, so that $q_S(x, t) \sim e^{-\kappa|x-x^*(t)| + \nu_S \kappa^2 t}$. Taking the logarithm, the free energy should have the form $F_{gs} \sim \kappa|x-x^*| - \nu_S \kappa^2 t$.

We now present a heuristic picture of how this behavior is modified for $\sigma_T > 0$, which allows us to quantitatively describe the saturation of the mean-squared error in Figs. 3 and 4. We compare the free-energy profiles for the $\sigma_T = 0$ and $\epsilon_T = 0$ cases. First, for $\sigma_T = 0$ and $\epsilon_T > 0$, F_{gs} discussed above cannot be the solution for early times, since if we were to start off with a flat initial condition, we would have $F(x, 0) = \text{cst}$, but the long-time solution above has $F_{gs}(\pm\infty, t) = \infty$. We ignore the motion of x^* , setting $x^* = 0$, and will justify this assumption afterwards. To look at the early-time transient behavior, we first relate Eq. (32) to a




FIG. 7. Schematic of the competition between the free energy of the planting $F_{\text{plant}} \sim t$ (left term) vs the KPZ free energy $F_{\text{KPZ}} \sim t^{2/3}$ (right term). At early times F_{KPZ} beats F_{plant} while eventually F_{plant} always wins.

Burgers-like equation for $u(x, t) = 2\nu_S \partial_x F(x, t)$,

$$\partial_t u + u \partial_x u = \nu_S \partial_x^2 u - 4\nu_S^2 \kappa \delta'(x). \quad (33)$$

The Burgers equation [i.e., without the $\delta'(x)$ on the right-hand side] has right- and left-moving soliton solutions [15] of the form $u_{\text{sol}}(x - vt)$ and $-u_{\text{sol}}(-x - vt)$, where

$$u_{\text{sol}}(z) = \frac{2v}{1 + \exp[(v/\nu_S)z]}. \quad (34)$$

We therefore make the ansatz, which we expect to be valid at long times:

$$u_{\text{plant}}(x, t) = \begin{cases} u_{\text{sol}}(x - vt) & x \geq 0 \\ -u_{\text{sol}}(-x - vt) & x < 0. \end{cases} \quad (35)$$

This ansatz has a jump at the origin that approaches $4v$ at long times. The $\delta'(x)$ term in Eq. (33) fixes $v = \nu_S \kappa$. The free energy is given by $F(x, t) = \frac{1}{2\nu_S} \int^x u(x', t) dx' + c(t)$, with the constant of integration $c(t)$ chosen so that $F(\pm\infty, t) = 0$. In the limit that the width of the solitons can be neglected, $F(x, t)$ takes the form

$$F_{\text{plant}}(x, t) = \begin{cases} -\nu_S \kappa^2 t + \kappa|x|, & |x| \leq \nu_S \kappa t \\ 0, & |x| > \nu_S \kappa t. \end{cases} \quad (36)$$

This is a "wedge" shape that is minimized at $x = 0$ that "carves" out downwards in time, shown on the left of Fig. 7. Our approximation that $x^* = 0$ is seen to be justified at long times because of the ballistic spreading of the wedge, which overwhelms the diffusive motion of x^* . Next, we look at the known scaling of the strong-coupling KPZ fixed point of the unplanted case $\epsilon_T = 0$, $\sigma_T > 0$. Reference [16] finds the exact scaling of the "height field" $h(x, t) := -F(x, t)$ again for the flat initial condition. The height field evolves as $\partial_t h = \nu \partial_x^2 h + \frac{\lambda_0}{2} (\partial_x h)^2 + \sqrt{D} \eta$. From this we can read off the mapping $\nu = \nu_S$, $\lambda_0/2 = \nu_S$, $\sqrt{D} = \epsilon_S \sigma_T / \sigma_S^2$. They show that the height field has an asymptotic form $\frac{\lambda_0}{2\nu} h_x = v_0 t + \lambda \xi_{x,t}$, where $\xi_{x,t} \sim O(1)$ is a random field in space at each time, as sketched on the right of Fig. 7; $\lambda = (1/2)(\bar{c}^2 t / T^5)^{1/3}$, $\bar{c} = D \lambda_0^2$, and $T = 2\nu$. v_0 is a constant background and therefore can be ignored. Converting from their variables to ours, we find that

$$F_{\text{KPZ}}(x, t) = 2^{-4/3} \epsilon_S^{4/3} \sigma_T^{4/3} \sigma_S^{-8/3} \nu_S^{-1/3} t^{1/3} \xi_{x,t}. \quad (37)$$

We now interpret the saturation of the mean-squared error as the result of competition between $F_{\text{plant}}(x, t)$ and $F_{\text{KPZ}}(x, t)$, as illustrated schematically in Fig. 7. Due to the t vs $t^{2/3}$ scaling,

we expect the KPZ effect to win until the planting effect is stronger, therefore when $F_{\text{plant}}(0, t_{\text{plateau}}) \sim F_{\text{KPZ}}(0, t_{\text{plateau}})$. This leads to the estimate

$$t_{\text{plateau}} \sim \frac{\sigma_S^2 \sigma_T^2 \nu_S}{\epsilon_S \epsilon_T^3}. \quad (38)$$

From Fig. 3, we have the picture that the student's walker grows as $x \sim t^{2/3}$ until we hit t_{plateau} , where it plateaus at $\overline{\text{RMSE}}(\infty) = x_{\text{plateau}} \sim t_{\text{plateau}}^{2/3}$. Substituting this in, we find that

$$\overline{\text{RMSE}}(\infty) \sim \frac{\sigma_S^{4/3} \sigma_T^{4/3} \nu_S^{2/3}}{\epsilon_S^{2/3} \epsilon_T^2}, \quad (39)$$

from which we predict that $x_{\text{plateau}} \sim \epsilon_T^{-2}$. This is consistent with the scalings we see in Fig. 4. The above picture supports an absence of a phase transition for one dimension, as the scaling holds for any parameters as long as they are not zero. This is because for the unplanted case $\epsilon_T = 0$, for any $\sigma_T > 0$, the system is always in the strong-coupling KPZ fixed point, which in turn is a consequence of the recurrence of random walks in one dimension. This is not true in the three-dimensional case, where there is a transition between the weak and strong disorder phases. Therefore, we may expect a transition in that case. We also note that the case of the delta initial condition would require a more involved ansatz but we expect the results to be similar due to the same underlying scaling.

In the planted ensemble picture, we can consider the posterior for path $x_{1:t}$ given $\phi_{1:t}$. Following the same procedure to go to the continuum, observables $O[x, x^*]$ dependent on entire paths can be formally written as

$$\begin{aligned} \mathbb{E}_{x, \eta, x^*} O[x, x^*] &= \int D\pi[\eta] \int_{p_T[x^*(0)]} D[x^*] e^{-E_0[x^*]} \\ &\times \frac{\int_{p_S[x(0)]} D[x] O[x, x^*] e^{-E[x, \eta, x^*]}}{Z_P(t)}. \end{aligned} \quad (40)$$

The case $\sigma_T = 0$ with the delta initial condition is the short-range interaction limit for the "memory model" considered in Ref. [17], but with open boundary conditions at the end times. This case corresponds to the somewhat unrealistic scenario where the student does not know that the images given by the teacher have no noise, and assumes that there is some level of noise. Reference [17] considers the long-time path-averaged mean-squared distance between the student and teacher paths, $\overline{\text{RMSE}}(\infty)$. Translating between the two languages, $1/4\nu_S = \beta/2$, $\beta u = \epsilon_S \epsilon_T / \sigma_S^2$, and $1/4\nu_T = \gamma/2$, we find

$$\overline{\text{RMSE}}_{\sigma_T=0}(\infty) \sim \frac{\sigma_S^2}{\epsilon_S \epsilon_T} (\nu_S + \nu_T). \quad (41)$$

For the delta initial condition, a quantity of interest is the long-time wandering exponent α that describes the long-time growth of fluctuations $\sqrt{\mathbb{E}[x_t^2]} \sim t^\alpha$. In the case $\epsilon_T = 0$, we have the directed polymer and have $\alpha = 2/3$. In case $\epsilon_S = 0$, looking at Eq. (31), we have the heat equation for q_S and have $\alpha = 1/2$. In the Bayesian optimal case, $S = T$, since the joint distribution is symmetric with $x_{1:t}$ and $x_{1:t}^*$, and because the teacher $x_{1:t}^*$ undergoes a random walk, we must also have $\alpha =$




FIG. 8. Plot of root-mean-squared width of the inferred path $\sqrt{\mathbb{E}[x_t^2]}$. At high teacher signal strength ϵ_T we see that the RMS width switches from $\sim t^{2/3}$ to $\sim t^{1/2}$ behavior at plateau time. As ϵ_T is lowered we see a crossover to $\sim t^{2/3}$ behavior at finite time.

$1/2$. What happens away from these limits is not immediately obvious.

Let us argue that $\alpha = 1/2$ if $\overline{\text{RMSE}}(\infty)$ is finite for a fixed ϵ_T . Since the RMS distance between the true and inferred paths is bounded, $x_t \sim x_t^* + C$ with $|C| \lesssim \overline{\text{RMSE}}(\infty)$. Since x_t^* can grow arbitrarily with time, at long times the x_t^* term dominates and fluctuations of x_t should grow as x_t^* does. In Fig. 8 we plot the RMS width of the inferred position, $\sqrt{\mathbb{E}[x_t^2]}$. We see that for large signal strength $\epsilon_T \sim 1$, the RMS width switches from $\sim t^{2/3}$ to $\sim t^{1/2}$ at some time t_{plateau} . Furthermore as ϵ_T is lowered we see a crossover from $\alpha = 1/2$ to $\alpha = 2/3$ for a given finite t .

IV. TREE CASE

We now consider the case where the random walker moves in a Cayley tree of depth t with branching number k , moving deeper into the tree at each time step, as shown by the location of the yellow nodes in Fig. 9. The Cayley tree can be seen as a Bethe approximation to a lattice as spatial dimension goes to infinity [18]. On the tree, positions can be labeled by which direction the walker took at each time step. Therefore we use the string notation $r_{1:t} = r_1 r_2 \dots r_t$ to denote positions on the tree, where each $r_\tau = 1, \dots, k$ denotes the direction chosen at depth τ . On a tree, there is only one path for a given position. Therefore we also can interpret $r_{1:t}$ as the path taken. If we assume that all branches are equally likely, i.e.,

$$p(r'_{1:t} | r_{1:t-1}) = \begin{cases} \frac{1}{k} & \text{if } r'_{1:t-1} = r_{1:t-1} \\ 0 & \text{otherwise,} \end{cases} \quad (42)$$

then the posterior can be written as

$$p_S(r_{1:t} | \phi_{1:t}) = \frac{e^{\epsilon_S \Phi(r_{1:t}) / \sigma_S^2}}{Z_S(\Phi)}, \quad (43)$$

where we define $\Phi(r_{1:t}) := \phi_{r_{1,1}} + \phi_{r_{1,2},2} + \dots + \phi_{r_{1,t},t}$, and $Z_S(\Phi) = \sum_{r_{1:t}} e^{\epsilon_S \Phi(r_{1:t}) / \sigma_S^2}$. In Fig. 9, the edges of the main figure are an example of the posterior $p_S(r_{1:t} | \phi_{1:t})$, inferred from the noisy images $\phi_{1:t}$ that are represented by the edges of the inset.




FIG. 9. An example of Bayesian inference on a random walker on a tree with branching number $k = 3$, $t = 4$, teacher and student signal strengths $\epsilon_T = \epsilon_S = 2$, and $\sigma_T = \sigma_S = 1$. The yellow nodes in the main figure are the true positions of the random walker while the color of each edge corresponds to the posterior for the path ending at the next node. Inset: The "noisy images" taken at each time step.

For a given inferred path $r_{1:t}$, let $w(r_{1:t}, r_{1:t}^*)$ be the fraction of time which $r_{1:t}$ is equal to the true path, $r_{1:t}^*$. Since on the tree, each path is equivalent to every other path, we can fix $r_{1:t}^*$ when calculating the overlap. The overlap is then

$$\begin{aligned} \bar{Y} &= \mathbb{E}_{r_{1:t}, \Phi} [w(r_{1:t}, r_{1:t}^*)] \\ &= \mathbb{E}_{\Phi} \left[\frac{\sum_{r_{1:t}} w(r_{1:t}, r_{1:t}^*) e^{\epsilon_S \Phi(r_{1:t}) / \sigma_S^2}}{Z_S(\Phi)} \right]. \end{aligned} \quad (44)$$

When there is no signal $\epsilon_T = 0$, we expect all k^t paths to be equally likely. Out of these paths, for a given overlap wt , there are k^{t-1-wt} paths. Then the average overlap is a geometric sum,

$$\bar{Y}_{\text{unif}}(t) = \frac{1}{t} \sum_{wt=0}^t wt k^{-1-wt} \xrightarrow{t \gg 1} \frac{1}{t \ln k}. \quad (45)$$

Following the prescription in Eq. (24), we separate the planted potential into the i.i.d. Gaussian part and the signal part as $\phi_{x,t} = \psi_{x,t} + \epsilon_T \delta_{x,x_t^*}$. Then, a path $r_{1:t}$ that has overlap q with the true path $r_{1:t}^*$ has the energy $\Psi(r_{1:t}) + \epsilon_T qt$, where $\Psi_{1:t}(r_{1:t}) = \psi_{r_{1,1}} + \psi_{r_{1,2},2} + \dots + \psi_{r_{1,t},t}$ is the unplanted random field. Therefore we can write Eq. (44) as an expectation on Ψ as

$$\bar{Y} = \mathbb{E}_{\Psi} \left[\frac{\sum_w w z_w(\Psi, r^*) e^{\epsilon_S \Psi(r_{1:t}) / \sigma_S^2}}{Z_P(\epsilon_T, \Psi)} \right], \quad (46)$$

where we introduced the partial partition function

$$z_w(\Psi, r^*) = \sum_{r_{1:t} \text{ with } w(r_{1:t}, r_{1:t}^*)=w} e^{\epsilon_S \Psi(r_{1:t}) / \sigma_S^2}, \quad (47)$$

and the total planted partition function now can be written as

$$Z_P(\epsilon_T, \Psi) = \sum_w z_w(\Psi, r^*) e^{\epsilon_S \Psi(r_{1:t}) / \sigma_S^2}. \quad (48)$$

Then $Z_P(0, \Psi)$ is the partition function for the directed polymer on a tree.

The directed polymer on a tree has been extensively studied, using the following three methods: the replica method, analogy to traveling waves [13,19], and as a special case of the so-called generalized random energy model (GREM) [20]. Below, we calculate the overlap between the true and inferred paths using the traveling wave and GREM approaches, and show that they lead to identical results.

A. GREM approach

The GREM was developed as an extension to the random energy model (REM), the simplest model for spin-glass systems, to include correlations between states.

Starting from the spin-glass language, in the GREM, each Ising spin configuration $\sigma = (\sigma_1, \dots, \sigma_N)$ is assigned a path $r_{1:t}$ on a tree with total number of paths $k_1 \dots k_t = 2^N$, with branching number k_i generally varying with depth i . Then, each bond on the tree is assigned an i.i.d. Gaussian random energy with width σ_i^2 again varying with depth. Then, each spin configuration is given random but correlated energies corresponding to summing up energies of bonds for that path.

Assigning the ferromagnetic configuration $\sigma = (+1, \dots, +1)$ to the "true" path $r_{1:t}^*$, the overlap with the ferromagnetic configuration can loosely be interpreted as the magnetization (per spin). Then, applying a magnetic field b shifts the energies of the configurations with overlap w by wbN .

In the case in which branching numbers and widths are constant with depth, the equations for the overlap with the true path from Bayesian inference and the magnetization due to a magnetic field coincide and are given by Eq. (46), with inverse temperature $\beta := \epsilon_S / \sigma_S^2$ and magnetic field $b = \epsilon_T / N$.

The case of the GREM in a magnetic field was studied by Derrida and Gardner in Ref. [20]. There they argue that in the large N limit, the quenched free energy can be approximated via "maximum *a posteriori*" approximation. In our case, where we are interested in the large t limit, this corresponds to

$$\lim_{t \rightarrow \infty} f = \max_w f_w, \quad (49)$$

where the quenched free energy (density) is $f := \frac{1}{t} \mathbb{E}_{\Psi} [\ln Z_P(\Psi)]$ and the partial free energy is

$$f_w := \frac{1}{t} \mathbb{E}_{\Psi} [\ln z_w(\Psi, r_{1:t}^*)] + \beta w \epsilon_T. \quad (50)$$

The GREM undergoes phase transitions, the structure of which is given by the variation in k_i and σ_i^2 . In the constant $k_i = k$ and $\sigma_i^2 = \sigma^2$ case, the GREM undergoes a single-phase transition like the REM.

Due to the structure of the tree, the partial partition function is another GREM [21]. The constant width and branching number case in the large t limit can be solved [13] to be

$$\frac{1}{t} \mathbb{E}_{\Psi} [\ln z_w(\Psi, r_{1:t}^*)] = (1 - w) \beta c \beta, \quad (51)$$

where the "speed of the front" c_β , whose naming will be made obvious later, is given by

$$c_\beta = \begin{cases} c(\beta) & \text{if } \beta \leq \beta_c \\ c(\beta_c) & \text{if } \beta > \beta_c, \end{cases} \quad (52)$$

with

$$c(\beta) = \frac{1}{\beta} \ln \left(k \int d\psi \rho_T(\psi) e^{-\beta\psi} \right), \quad (53)$$

$\rho_T(\psi) = \pi_T(\psi)$ is the teacher's distribution of the i.i.d. noise, and the critical inverse temperature β_c is given by $\frac{\partial}{\partial \beta} c(\beta)|_{\beta=\beta_c} = 0$. The w dependence of Eq. (50) can be found to be

$$f_w = \begin{cases} (\beta\epsilon_T - [\frac{\ln k}{\beta} + \frac{\beta\sigma_T^2}{2}])w + C_1 & \text{if } \beta \leq \beta_c \\ (\beta\epsilon_T - \sigma_T\sqrt{2\ln k})w + C_2 & \text{if } \beta > \beta_c, \end{cases} \quad (54)$$

where C_1, C_2 are constants independent of w and $\beta_c = \sqrt{2\ln k}/\sigma_T$. As f_w is linear in w , we see that w that dominates the free energy is either 0 or 1, determined by the sign of the gradient term. The line where the gradient is zero then gives us the boundary between the two regimes, and can be found to be

$$\frac{\epsilon_T}{\sigma_T} = \begin{cases} \ln k \left(\frac{\epsilon_S \sigma_T}{\sigma_S^2} \right)^{-1} + \frac{1}{2} \frac{\epsilon_S \sigma_T}{\sigma_S^2} & \text{if } \frac{\epsilon_S \sigma_T}{\sigma_S^2} \leq \sqrt{2\ln k} \\ \sqrt{2\ln k} & \text{if } \frac{\epsilon_S \sigma_T}{\sigma_S^2} > \sqrt{2\ln k}. \end{cases} \quad (55)$$

B. Traveling waves approach

The overlap can also be calculated using the traveling waves approach. To do this we extend the arguments of [13] to the planted case. Using the structure of the tree, up to a multiplicative constant, the recursion relation for planted partition function can be written as

$$Z_P(\epsilon_T, t+1) = e^{\beta(\psi - \epsilon_T)} \left(Z_P^{(1)}(\epsilon_T, t) + \sum_{i=2}^k Z_P^{(i)}(0, t) \right), \quad (56)$$

with the initial condition imposed to be $Z_P(\epsilon_T, 0) = 1$. We note that $\epsilon_T = 0$ retrieves the unplanted partition function of the directed polymer. Defining the planted generating function as

$$G_{\epsilon_T}(x, t) := \mathbb{E}_\Psi \{ \exp[-Z_P(\epsilon_T, t) e^{-\beta x}] \}, \quad (57)$$

the recursion relation of the planted generating function can be written as

$$G_{\epsilon_T}(x, t+1) = \mathbb{E}_\Psi \left[G_{\epsilon_T}(x + \psi - \epsilon_T, t) \times G_0(x + \psi - \epsilon_T, t)^{k-1} \right], \quad (58)$$

where $\psi \sim \mathcal{N}(0, \sigma_T^2)$ represents noise at a single site or, equivalently, bond.

Let us consider continuous time evolution, where in a time step δt , we can have branching with probability $\lambda \delta t$ (so $k = 2$), and no branching with probability $(1 - \lambda \delta t)$. We renormalize the noise to be $\psi \sim \mathcal{N}(0, 2D\delta t)$ and the signal to be $\epsilon_T \delta t \rightarrow \epsilon_T$. Then G_{ϵ_T} evolves as

$$\partial_t G_{\epsilon_T} = D \partial_x^2 G_{\epsilon_T} - \epsilon_T \partial_x G_{\epsilon_T} - \lambda(1 - G_0)G_{\epsilon_T}. \quad (59)$$




FIG. 10. Numerical simulations of planted and unplanted generating functions. At $t = 0$, the unplanted (G_0) and planted (G_{ϵ_T}) generating functions are equal (gray dashed line). The characteristic velocity of G_{ϵ_T} is ϵ_T , while G_0 travels with c_β . Top: $\epsilon_T = c_\beta/2 < c_\beta$ for $t = 60$. G_{ϵ_T} (orange) sticks to G_0 (blue). Bottom: As top but for $\epsilon_T = 3c_\beta/2 > c_\beta$. Here we see that G_{ϵ_T} detaches from G_0 , and moves with its characteristic velocity ϵ_T .

There are three components to Eq. (59). The first term proportional to D is the diffusive term. The second term moves the planted generating function with characteristic velocity ϵ_T towards the right. The third term is like the nonlinear term in the Fisher-Kolmogorov-Petrovsky-Piskunov (FKPP) equation that pushes down $G_{\epsilon_T}(x, t)$ at positions x where $G_0(x, t) < 1$.

It is known that the unplanted generating function $G_0(x, t)$ evolves according to the FKPP equation. In the long-time limit, G_0 moves as a profile that switches from 0 to 1, and moves with speed c_β given by Eq. (52), $\lim_{t \rightarrow \infty} G_0(x, t) \approx w_\beta(x - c_\beta t)$, with $w_\beta(-\infty) = 0$ and $w_\beta(\infty) = 1$.

G_0 and G_{ϵ_T} start with the same initial conditions. If $\epsilon_T < c_\beta$, then the third term determines the position of the front and G_{ϵ_T} moves with the same velocity as G_0 . If $\epsilon_T > c_\beta$, then the support of G_{ϵ_T} moves to a region where $G_0 = 1$ and therefore the third term vanishes, leading to a linear partial differential equation for an error function that moves towards the right with speed ϵ_T and spreads as $\sqrt{t}$. Therefore G_{ϵ_T} moves with velocity $v_{\beta, \epsilon_T} = \max(\epsilon_T, c_\beta)$. This is supported and illustrated with numerical simulations in Figs. 10 and 11.

Now $G_{\epsilon_T}(x, t)$ switches from 0 to 1 approximately at point $\beta \hat{x} \sim \ln Z_P(\epsilon_T, t)$. From the previous arguments, this point at long times is $\hat{x}(t) = v_{\beta, \epsilon_T} t$. Assuming that fluctuations in $\ln Z_P(\epsilon_T, t)$ are sublinear, this means that $\mathbb{E}_\Psi[\ln Z_P(\epsilon_T, t)] = \beta v_{\beta, \epsilon_T} t$.

From the definition of average overlap Y , we have that $Y = \frac{1}{\beta t} \frac{\partial}{\partial \epsilon_T} \mathbb{E}_\Psi[\ln Z_P(\epsilon_T, t)] = \frac{\partial}{\partial \epsilon_T} v_{\beta, \epsilon_T}$. Then, when $\epsilon_T \leq c_\beta$, v_{β, ϵ_T} does not change with ϵ_T and therefore $Y = 0$. On the other




FIG. 11. Heatmap: Unplanted generating function (G_0) evolving according to the FKPP equation. Blue dashed line: The front of the planted generating function (G_{ϵ_T}), initially moving with characteristic velocity $\epsilon_T = c_\beta/2$ (red dashed line) smaller than the unplanted velocity c_β . The front of G_{ϵ_T} eventually sticks to that of G_0 . Green dashed line: For characteristic velocities larger than c_β ($\epsilon_T = 3c_\beta/2 > c_\beta$), G_{ϵ_T} is free to travel at ϵ_T .

hand, when $\epsilon_T > c_\beta$, then we have $\partial_{\epsilon_T} v_{\beta, c_{\epsilon_T}} = 1$ and therefore have the overlap of $Y = 1$. As c_β has a phase transition, we must consider the cases when $\beta \leq \beta_c$ and $\beta > \beta_c$ after which we retrieve the identical phase boundary, Eq. (55).

C. General noise distribution

We can generalize to the case where the teacher uses an arbitrary noise distribution, whilst the student conducts inference assuming Gaussian noise. This can be done by choosing a different distribution for $\rho_T(\psi)$ in Eq. (53). The minimum speed is again given by $\partial_\beta c(\beta)|_{\beta=\beta_c} = 0$. For example, for the top-hat distribution $\rho_T(\psi) = 1/2a$ for $-a < \psi < a$, we have $c(\beta) = \ln[k \sinh(\beta a)/\beta a]/\beta$, with minimum given at $\beta_c = x(k)/a$ for $x \coth(x) - \ln[k \sinh(x)/x] = 1$. This is the same qualitative picture as Fig. 12, where c_β decreases until it hits minimum at some β_c and then freezes. Numerically testing on various unimodal distributions, we consistently found the same qualitative picture.

D. Phase diagram and numerics

The phase diagram obtained via either approach with the boundary given by Eq. (55) is sketched in Fig. 12.

In the Bayes optimal case, corresponding to the 45° line in Fig. 12, the planted directed polymer "inherits" the unplanted model's transition point, $\epsilon/\sigma = \sqrt{2 \ln k}$. Additionally, for a given ϵ_T/σ_T , one cannot get a greater overlap than the Bayes optimal case. This is expected as it is a standard result that one cannot do better than the Bayes optimal case for various tasks such as classification and estimation [22,23]. However, for the specific case of the tree, it appears that as long as $\epsilon_S > \sqrt{2 \ln k}$, performance in inference is only limited by ϵ_T . This suggests that in the case where the student does not know




FIG. 12. Phase diagram for overlap Y on the plane of teacher-student signal strengths $\epsilon_S - \epsilon_T$ for the case of a tree with branching number k . $Y = 1$ corresponds to perfect inference whilst $Y = 0$ corresponds to failure of inference. The boundary corresponds to a first-order phase transition.

the true value of the signal strength, they can pick any value $\epsilon_S > \sqrt{2 \ln k}$ and will be able to infer as well as the Bayes optimal case. A separate partially unanswered question is how sure the student is about their answer; that is, the variance or uncertainty on their inferred path. This could be quantified by the overlap between two paths $r_{1:t}^{(1)}, r_{1:t}^{(2)}$ sampled from the student's posterior,

$$\bar{Y}_S := \frac{1}{t} \mathbb{E}_{y_{1:t}} \left\{ \mathbb{E}_{r_{1:t}^{(1)}, r_{1:t}^{(2)} \sim p_S(\cdot | y_{1:t})} \left[\sum_{\tau=1}^t \delta_{r_\tau^{(1)}, r_\tau^{(2)}} \right] \right\}. \quad (60)$$

The case of zero teacher signal-to-noise ratio ($\epsilon^* = 0$) corresponding to the usual directed polymer case was studied in [13], which revealed a transition from $\bar{Y}_S = 0$ to 1 at $\epsilon_S\sigma_T/\sigma_S^2 = \sqrt{2 \ln k}$, coincides with our transition point at Bayes optimality. The symmetry between the student and teacher means that $\bar{Y}_S = \bar{Y}$ in the Bayes optimal case, which immediately gives us $\bar{Y}_S$ along the line $\epsilon_T/\sigma_T = \epsilon_S\sigma_T/\sigma_S^2$ in Fig. 12. This agrees with the transition in the averaged entropy of the posterior, in the Bayes optimal case found in the recent preprint, Ref. [24], for a generalized measurement model. It would be interesting to extend either of these quantities to the rest of the phase diagram.

The reason for the first-order "all-to-nothing" transition in the tree case is due to the fact that in the thermodynamic limit all overlap (w) dependent contributions to the free energy are linear. This is because any path with overlap w is a combination of the path following the planted trajectory (with purely energetic and no entropic contribution to the free energy) followed by a path which has left it, never to return, and is therefore accounted for by the solution of the unplanted problem. The transition between $w = 0$ and $w = 1$ is therefore determined purely by the sign of the linear term. This is not the case in the finite-dimensional problem on a lattice, for example.

To confirm our theory, in Fig. 13, we numerically study the overlap as we vary time t up to $t = 11$ for fixed $\sigma_S = \sigma_T = \sigma = 1$, $k = 2$, and $\epsilon_S = 4\sigma\sqrt{2 \ln k}$ for $\#\text{disorder} = 1000$ disorder realizations. We confirm the "finite-size scaling" of




FIG. 13. Numerical study of the average overlap Y with teacher signal strength ϵ_T at various times t at fixed branching number k , student signal strength, and number of disorder realisations $\#_{\text{disorder}}$.

Eq. (45) for $\epsilon_T = 0$, and see that the numerics support theoretical transition point $\epsilon_T^* = \sigma \sqrt{2 \ln k}$.

V. CONCLUSION

In this work we introduced the planted directed polymer problem. Our main contributions are as follows: (1) We derived an explicit mapping of the inference problem to the directed polymer. (2) We made an explicit connection to the statistical physics of inference in the planted ensemble. (3) We solved the planted polymer on a tree exactly, adapting existing machinery developed for the unplanted case. (4) We provide heuristic arguments combining numerical and analytical reasoning for the absence of a phase transition in the one-dimensional problem, as well as discussing qualitative features of the free-energy profile and the higher-dimensional versions of the problem.

The earlier work of Yuille and Coughlan [7], studied tracking of roads in aerial images as a problem of Bayesian inference, including the notion of the planted ensemble. They raised the question of whether there is a fundamental limit to the success of inference in this situation and connected this

limit to a phase transition in the thermodynamic limit. They made the further simplification of restricting the paths to a tree, as we did in Sec IV, but without connecting to the earlier literature on the directed polymer. They obtain a bound on the "error" $t(1 - Y)$ that depends on a regularization parameter used to treat the $t \rightarrow \infty$ limit but not on t . Therefore the regime of vanishing overlap $Y = 0$ is not handled correctly. In Ref. [25] the same authors explored the effect of nonoptimality due to inference based on the wrong prior.

The More recent work of Offer [8] addressed the related problem of Bayesian tracking in *clutter* using a mean-field approximation as well as numerical evaluation of the posterior, finding a phase transition in three or more spatial dimensions. Again, the connection to the directed polymer was not made.

A separate but related question is whether or not a random walker exists in the images at all. References [26,27] addressed this question by testing the hypothesis of there being a random walker versus the null hypothesis of no walker (i.e., images of white noise). They relate the decay rate of the false-negative probability to the average log-likelihood ratio between the hypothesis and the null hypothesis. The log-likelihood ratio coincides with the free energy of the directed polymer. It would be interesting to apply the known results of the directed polymer literature to this setting.

Several other avenues for future research exist. One could attempt to adapt existing analytical tools for the $d = 1$ directed polymer (such as the replica Bethe ansatz; see, e.g., Ref. [28]) to the planted case in order to calculate some of the observables explored in this paper. Moving to higher dimensions, it is known that the directed polymer has a transition at finite disorder strength for $d \geq 3$. We may expect the planted directed polymer to "inherit" the phase transition, as suggested by the investigation of Ref. [8].

ACKNOWLEDGMENTS

S.W.P.K. and A.L. acknowledge support from EPSRC DTP International Studentship Grant Ref. No. EP/W524475/1 and EPSRC Critical Mass Grant No. EP/V062654/1, respectively. We are grateful to Bernard Derrida for useful discussions.

- [1] L. Zdeborová and F. Krzakala, *Adv. Phys.* **65**, 453 (2016).
- [2] Note that throughout we adopt the convention that different distributions are distinguished by the names of their arguments: a convenient abuse of notation.
- [3] L. Rabiner and B. Juang, *IEEE ASSP Mag.* **3**, 4 (1986).
- [4] L. R. Rabiner, *Proc. IEEE* **77**, 257 (1989).
- [5] G. Welch and G. Bishop, An introduction to the Kalman filter, Technical Report (University of North Carolina at Chapel Hill, Chapel Hill, NC, 1995).
- [6] I. Smal, W. Niessen, and E. Meijering, in *Third IEEE International Symposium on Biomedical Imaging: Nano to Macro* (IEEE, New York, 2006), pp. 550–553.
- [7] A. L. Yuille and J. M. Coughlan, *IEEE Trans. Pattern Anal. Mach. Intell.* **22**, 160 (2000).

- [8] C. R. Offer, in *2018 21st International Conference on Information Fusion (FUSION)* (IEEE, New York, 2018), pp. 2129–2136.
- [9] M. Kardar, *Statistical Physics of Fields* (Cambridge University Press, Cambridge, 2007).
- [10] B. Skinner, J. Ruhman, and A. Nahum, *Phys. Rev. X* **9**, 031009 (2019).
- [11] Y. Li, X. Chen, and M. P. A. Fisher, *Phys. Rev. B* **100**, 134306 (2019).
- [12] A. C. Potter and R. Vasseur, *Entanglement in Spin Chains: From Theory to Quantum Technology Applications* (Springer, New York, 2022), pp. 211–249.
- [13] B. Derrida and H. Spohn, *J. Stat. Phys.* **51**, 817 (1988).

- [14] R. L. Stratonovich, *Non-Linear Transformations of Stochastic Processes* (Elsevier, New York, 1965), pp. 427–453.
- [15] P. J. Olver *et al.*, *Introduction to Partial Differential Equations* (Springer, New York, 2014), Vol. 1, p. 317.
- [16] P. Calabrese and P. Le Doussal, *Phys. Rev. Lett.* **106**, 250603 (2011).
- [17] V. Dotsenko, *J. Stat. Mech.* (2022) 093302.
- [18] S. Katsura and M. Takizawa, *Prog. Theor. Phys.* **51**, 82 (1974).
- [19] B. Derrida and E. Gardner, *J. Phys. C* **19**, 2253 (1986).
- [20] B. Derrida and E. Gardner, *J. Phys. C* **19**, 5783 (1986).
- [21] Note that the partial partition function will have a different branching number in the first time step, leading to a correction term compared to when there is equal branching numbers. However, this correction is subextensive.
- [22] C. M. Bishop and N. M. Nasrabadi, *Pattern Recognition and Machine Learning* (Springer, New York, 2006), Vol. 4.
- [23] H. Pishro-Nik, *Introduction to Probability, Statistics, and Random Processes* (Kappa Research, LLC Blue Bell, PA, 2014), Chap. 9.1.5.
- [24] F. Gerbino, G. Giachetti, P. Le Doussal, and A. De Luca, Measurement-induced phase transition in state estimation of chaotic systems and the directed polymer, [arXiv:2501.00547](https://arxiv.org/abs/2501.00547).
- [25] A. L. Yuille and J. Coughlan, in *Proceedings of the 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)* (IEEE, New York, 1999), Vol. 2, pp. 631–637.
- [26] M. Burnashev, *Theory Probab. Appl.* **26**, 554 (1982).
- [27] A. Agaskar and Y. M. Lu, [arXiv:1504.06924](https://arxiv.org/abs/1504.06924).
- [28] P. Le Doussal and P. Calabrese, *J. Stat. Mech.* (2012) P06001.