

New Method for Parameter Estimation in Probabilistic Models: Minimum Probability Flow

Jascha Sohl-Dickstein,^{1,4} Peter B. Battaglino,^{2,4} and Michael R. DeWeese^{2,3,4}

¹*Biophysics Graduate Group, Berkeley, 94720*

²*Department of Physics, Berkeley, 94720*

³*Helen Wills Neuroscience Institute, Berkeley, 94720*

⁴*Redwood Center for Theoretical Neuroscience University of California, Berkeley, 94720*

(Received 12 April 2011; revised manuscript received 17 April 2011; published 21 November 2011)

Fitting probabilistic models to data is often difficult, due to the general intractability of the partition function. We propose a new parameter fitting method, minimum probability flow (MPF), which is applicable to any parametric model. We demonstrate parameter estimation using MPF in two cases: a continuous state space model, and an Ising spin glass. In the latter case, MPF outperforms current techniques by at least an order of magnitude in convergence time with lower error in the recovered coupling parameters.

DOI: 10.1103/PhysRevLett.107.220601

PACS numbers: 05.20.-y, 02.50.Cw, 02.50.Ng, 75.10.Nr

Scientists and engineers increasingly confront large and complex data sets that defy traditional modeling and analysis techniques. For example, fitting well-established probabilistic models from physics to population neural activity recorded in retina [1,2] or cortex [3–5] is currently impractical for populations of more than about 100 neurons [6]. Similar difficulties occur in many other fields, including computer science [7], genomics [8], and physics [9].

The canonical problem is to find parameter values θ that result in the best match between a model and a list of observations \mathcal{D} . Parameter estimation can be viewed as the inverse of the usual problem physicists face: rather than assuming fixed model parameters, such as coupling strengths in an Ising spin glass, and then predicting observable properties of the system, such as spin-spin correlations, our goal is to start with a series of observations and then estimate the underlying model parameters. This is a challenging problem in many interesting cases [7,9,10].

Consider a data distribution over N discrete states, represented as a vector $\mathbf{p}^{(0)} \in \mathbb{R}^N$, with $p_i^{(0)}$ the fraction of the observations \mathcal{D} in state i . A model distribution parameterized by θ is similarly represented as $\mathbf{p}^{(\infty)}(\theta) \in \mathbb{R}^N$. The superscripts (0) and (∞) indicate initial conditions and equilibrium under system dynamics, as described below. For any model distribution $\mathbf{p}^{(\infty)}(\theta)$, the probability assigned to each state can be written

$$p_i^{(\infty)}(\theta) = \frac{\exp[-E_i(\theta)]}{Z(\theta)}, \quad (1)$$

where $\mathbf{E}(\theta) \in \mathbb{R}^N$ can be viewed as the energy in the familiar Boltzmann distribution (with $k_B T$ set to 1). $Z(\theta)$ is the partition function, which involves a sum over all N possible states of the system,

$$Z(\theta) = \sum_i^N \exp[-E_i(\theta)]. \quad (2)$$

For clarity we develop our method using discrete state spaces, but it extends to probabilistic models over continuous state spaces, as we demonstrate for an independent component analysis (ICA) model [11].

The standard objective for parameter estimation is to maximize the likelihood of the model $\mathbf{p}^{(\infty)}(\theta)$ given the observations \mathcal{D} , or equivalently to minimize $D_{\text{KL}}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(\infty)}(\theta))$, the Kullback-Leibler (KL) divergence between the data distribution and model distribution [7,12]. The estimated parameters $\hat{\theta}$ are given by

$$\hat{\theta} = \operatorname{argmin}_{\theta} D_{\text{KL}}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(\infty)}(\theta)), \quad (3)$$

$$D_{\text{KL}}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(\infty)}(\theta)) = \sum_i p_i^{(0)} \log p_i^{(0)} - \sum_i p_i^{(0)} \log p_i^{(\infty)}(\theta). \quad (4)$$

Unfortunately, the partition function $Z(\theta)$ in $p_i^{(\infty)}(\theta)$ usually involves an intractable sum over all system states. This is commonly the major impediment to parameter estimation.

Many approaches exist for approximate parameter estimation, including mean field theory [13,14] and its expansions [14,15], variational Bayes techniques [16,17], pseudolikelihood [18], contrastive divergence [19,20], score matching [21,22], minimum velocity learning [23], and a multitude of Monte Carlo and numerical integration-based methods [10,24].

Most Monte Carlo methods rely on two core concepts from statistical physics, which we will use to develop our parameter estimation technique, minimum probability flow (MPF). The first of these is conservation of probability, as enforced by the master equation for the evolution of a distribution $\mathbf{p}^{(t)}$ with time

$$\dot{p}_i^{(t)} = \sum_{j \neq i} \Gamma_{ij}(\theta) p_j^{(t)} - \sum_{j \neq i} \Gamma_{ji}(\theta) p_i^{(t)}. \quad (5)$$

$\Gamma_{ij}(\theta)$ gives the rate at which probability flows from state j into state i . The first term of Eq. (5) captures the flow of probability out of other states j into the state i , and the second represents the flow out of i into other states j .

The second core concept is detailed balance,

$$\Gamma_{ji}p_i^{(\infty)}(\theta) = \Gamma_{ij}p_j^{(\infty)}(\theta), \quad (6)$$

which when satisfied ensures that the probability flow from state i into state j equals the probability flow from j into i , and thus that the distribution $\mathbf{p}^{(\infty)}$ is a fixed point of the dynamics. Sampling in most Monte Carlo methods is performed by choosing Γ consistent with Eq. (6) (and the added requirement of ergodicity [7]), then stochastically running the dynamics in Eq. (5). Unfortunately, sampling algorithms can be exceedingly slow to converge, and are thus ill suited for use in each parameter update step during parameter estimation.

Using these two core concepts, we propose an approach, illustrated in Fig. 1, that avoids both sampling and explicit calculation of the partition function. Specifically, we establish deterministic dynamics obeying Eqs. (5) and (6) that converge to the model distribution $\mathbf{p}^{(\infty)}(\theta)$, and initialize them at the data distribution $\mathbf{p}^{(0)}$. Rather than allowing the dynamics to fully converge and making parameter updates that minimize $D_{\text{KL}}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(\infty)}(\theta))$ as in maximum likelihood learning [Eqs. (3) and (4)], our parameter updates instead minimize $D_{\text{KL}}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(\epsilon)}(\theta))$, the KL divergence after running the dynamics for an infinitesimal time ϵ . This requires computing only the instantaneous flow of probability away from the data distribution at time $t = 0$.




FIG. 1 (color). An illustration of parameter estimation using MPF. The three successive panels illustrate the sequence of parameter updates that occur during learning. Each set of axes represents the space of probability distributions. The dashed red curves indicate the family of model distributions $\mathbf{p}^{(\infty)}(\theta)$ parametrized by θ . The black solid curves indicate deterministic dynamics that transform the data distribution $\mathbf{p}^{(0)}$ into the model distribution $\mathbf{p}^{(\infty)}(\theta)$. Under maximum likelihood learning, model parameters θ are chosen so as to minimize the KL divergence between the data distribution $\mathbf{p}^{(0)}$ and the model distribution $\mathbf{p}^{(\infty)}(\theta)$. Under MPF the KL divergence between $\mathbf{p}^{(0)}$ and $\mathbf{p}^{(\epsilon)}(\theta)$ is minimized instead, where $\mathbf{p}^{(\epsilon)}(\theta)$ is the distribution obtained by initializing the dynamics at the data distribution $\mathbf{p}^{(0)}$ and then evolving them for an infinitesimal time ϵ . Here we represent graphically how parameter updates that pull $\mathbf{p}^{(\epsilon)}(\theta)$ towards $\mathbf{p}^{(0)}$ also tend to pull $\mathbf{p}^{(\infty)}(\theta)$ towards $\mathbf{p}^{(0)}$.

The transition rates Γ_{ij} are underconstrained by Eq. (6). Introducing the additional constraint that Γ be invariant to the addition of a constant to the energy function [as is true for the model distribution $\mathbf{p}^{(\infty)}(\theta)$], we choose the following form for Γ_{ij} :

$$\Gamma_{ij} = g_{ij} \exp\left[\frac{1}{2}[E_j(\theta) - E_i(\theta)]\right], \quad (i \neq j), \quad (7)$$

where $g_{ij} = g_{ji} \in \{0, 1\}$. The vast majority of the factors g_{ij} can generally be set to 0. However, for the dynamics to converge to $\mathbf{p}^{(\infty)}(\theta)$, there must be sufficient nonzero Γ elements to ensure mixing. In binary systems, good results are obtained by setting $g_{ij} = g_{ji} = 1$ only for states i and j differing by a single bit flip. The elements of g_{ij} may also be sampled, rather than set by a deterministic scheme (see Supplemental Material [25]).

Given the transition matrix Γ and the list \mathcal{D} of observed data samples, and taking ϵ small, the objective function $D_{\text{KL}}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(\epsilon)}(\theta))$ is approximated by its first order Taylor expansion, denoted $K(\theta)$ (see Supplemental Material [25])

$$K(\theta) = D_{\text{KL}}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(\epsilon)}(\theta))|_{t=0} + \epsilon \left. \frac{\partial D_{\text{KL}}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(\epsilon)}(\theta))}{\partial t} \right|_{t=0}, \quad (8)$$

$$= \epsilon \sum_{i \notin \mathcal{D}} \dot{p}_i^{(0)} = \frac{\epsilon}{M} \sum_{i \notin \mathcal{D}} \sum_{j \in \mathcal{D}} \Gamma_{ij}, \quad (9)$$

$$= \frac{\epsilon}{M} \sum_{j \in \mathcal{D}} \sum_{i \notin \mathcal{D}} g_{ij} \exp\left[\frac{1}{2}[E_j(\theta) - E_i(\theta)]\right], \quad (10)$$

where $M = |\mathcal{D}|$ is the number of data samples. Parameter estimation is performed by finding $\hat{\theta} = \text{argmin}_{\theta} K(\theta)$, generally via gradient descent of $K(\theta)$. Thus, minimizing the KL divergence $D_{\text{KL}}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(\epsilon)}(\theta))$ for small ϵ is equivalent to minimizing the initial flow of probability out of data states j into nondata states i [Eq. (9)]. For small systems, or large numbers of observations, every state may be a data state, in which case the first order term vanishes and higher order terms must be included.

The dimensionalities of $\mathbf{p}^{(0)}$ and Γ are typically large (e.g., 2^d and $2^d \times 2^d$, respectively, for a d -bit binary system). Fortunately, both can also be made extremely sparse: $p_j^{(0)} = 0$ for all nondata states $j \notin \mathcal{D}$, and we need only evaluate Γ_{ij} for which $j \in \mathcal{D}$ and $g_{ij} \neq 0$. The cost in both memory and time is therefore only $\mathcal{O}(M)$ per learning step. The dependence of total convergence time on the number of samples M is problem dependent, but it is roughly $\mathcal{O}(M)$ for the Ising spin glass model discussed below (see Supplemental Material [25]).

In addition, when estimating parameters for a model in the exponential family—that is, models such as spin glasses for which the energy function is linear in the parameters θ —the MPF objective function $K(\theta)$ is convex [26], guaranteeing that the global minimum can always be found via gradient descent. For exponential family models

over continuous rather than discrete state spaces, MPF further provides a closed form solution for parameter estimation [27]. MPF is also consistent—meaning that if the data distribution $\mathbf{p}^{(0)}$ belongs to the family of distributions $\mathbf{p}^{(\infty)}(\theta)$ parametrized by θ (Fig. 1, red dashed line), the objective function $D_{\text{KL}}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(\infty)}(\theta))$ will have its global minimum at the true parameter values (see Supplemental Material [25]).

We evaluated performance by fitting an Ising spin glass (sometimes referred to in the computer science literature as a fully visible Boltzmann machine or simply as an Ising model) of the form

$$p^{(\infty)}(\mathbf{x}; \mathbf{J}) = \frac{1}{Z(\mathbf{J})} \exp[-\mathbf{x}^T \mathbf{J} \mathbf{x}], \quad (11)$$

where \mathbf{J} only had nonzero elements corresponding to nearest-neighbor units in a two-dimensional square lattice, and bias terms along the diagonal. The training data \mathcal{D} consisted of M d -element independent and identically distributed (iid) binary samples $\mathbf{x} \in \{0, 1\}^d$ generated via Swendsen-Wang sampling [28] from a spin glass with known coupling parameters. In this example, we used a square 10×10 lattice, $d = 10^2$. The nondiagonal nearest-neighbor elements of \mathbf{J} were set using draws from a normal distribution with variance $\sigma^2 = 10$. The diagonal (bias) elements of \mathbf{J} were set so that each column of \mathbf{J} summed to 0, and the expected unit activations were 0.5. The $2^d \times 2^d$ element transition matrix Γ was populated sparsely,

$$g_{ij} = g_{ji} = \begin{cases} 1 & \text{states } i \text{ and } j \text{ differ by single bit flip} \\ 0 & \text{otherwise} \end{cases}. \quad (12)$$

Code implementing MPF is available [29].

We compared parameter estimation using MPF against parameter estimation using four competing techniques: mean field theory (MFT) with Thouless-Anderson-Palmer (TAP) corrections [30], one-step and ten-step contrastive divergence [19] (CD-1 and CD-10), and pseudolikelihood [18]. The results of our simulations are shown in Fig. 2, which plots the mean square error in the recovered \mathbf{J} and in the corresponding pairwise correlations as a function of learning time for MPF and the competing approaches outlined above. Using MPF, learning took approximately 60 seconds, compared to roughly 800 seconds for pseudolikelihood and approximately 20 000 seconds for both 1-step and 10-step contrastive divergence. Reasonable steps were taken to optimize the performance of all the parameter estimation methods tested (see Supplemental Material [25]). Note that, given sufficient samples, MPF is guaranteed to converge exactly to the right answer, as it is consistent and the objective function $K(\theta)$ is convex for a spin glass. MPF fit the model to the data more accurately than any of the other techniques. MPF was dramatically faster to converge than any of the other techniques tested, with the exception of MFT + TAP, which failed to find reasonable parameters. Note that MFT + TAP does converge to the

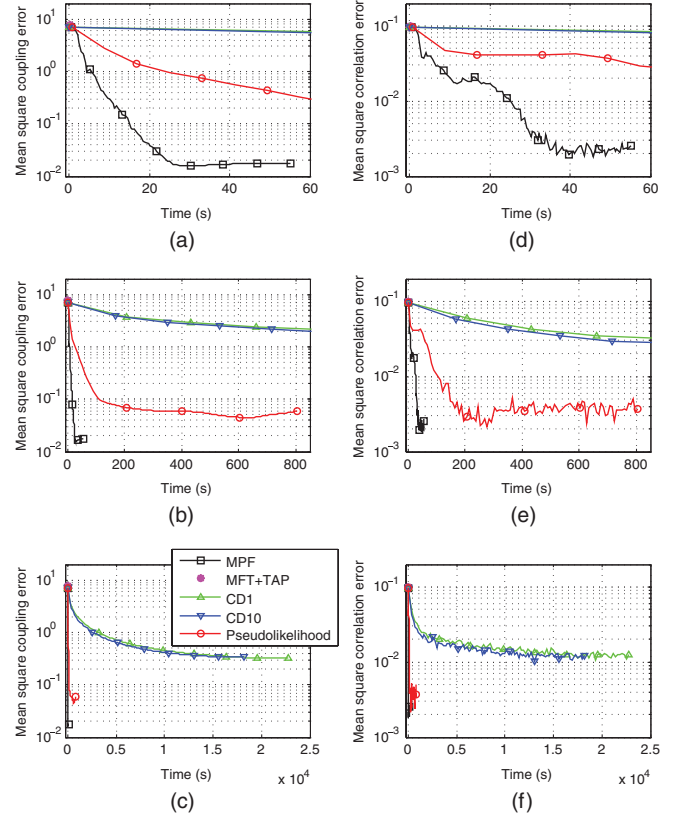


FIG. 2 (color). A demonstration of MPF outperforming existing techniques for parameter recovery in an Ising spin glass model. (a) Time evolution of the mean square error in the coupling strengths for 5 methods for the first 60 seconds of learning. Note that mean field theory with second order corrections (MFT + TAP) actually increases the error above random parameter assignment, though it does converge to the correct answer in some other parameter regimes, such as in the high temperature limit of this Ising spin glass model [15]. (b) Mean square error in the coupling strengths for the first 800 seconds of learning. (c) Mean square error in coupling strengths for the entire learning period. (d)–(f) Mean square error in pairwise correlations for the first 60 seconds of learning, the first 800 seconds of learning, and the entire learning period, respectively. In every comparison above MPF finds a better fit, and for all cases but MFT + TAP does so in a shorter time.

correct answer in certain parameter regimes, such as the high temperature limit [15], while remaining much faster than the other four techniques.

As a demonstration of parameter estimation using MPF for a continuous state space probabilistic model, we trained the filters $\mathbf{J} \in \mathbb{R}^{d \times d}$ of a d dimensional ICA [11] model with a Laplace prior,

$$p^{(\infty)}(\mathbf{x}; \mathbf{J}) = \frac{e^{-\sum_k |\mathbf{J}_k \mathbf{x}|}}{2^d \det(\mathbf{J}^{-1})}, \quad (13)$$

where $\mathbf{x} \in \mathbb{R}^d$ is a continuous state space. Since the log likelihood and its gradient can be calculated analytically




FIG. 3. A demonstration of MPF parameter estimation in a continuous state space probabilistic model. Each square represents a 10×10 pixel filter \mathbf{J}_k for an ICA model trained on natural image patches via (a) MPF or (b) maximum likelihood learning. The visual similarity of the filters is consistent with the nearly identical average log likelihoods for the two models (-174.0 bits and -173.6 bits, respectively).

for ICA, we solved for \mathbf{J} via maximum likelihood learning [Eq. (3)] as well as MPF, and compared the resulting log likelihoods. Training data consisted of natural image patches. The log likelihood of the model trained with MPF was -174.0 bits, while that for the maximum likelihood trained model was a nearly identical -173.6 bits. Average log likelihood at parameter initialization was -273.0 bits (see Supplemental Material [25]). The edge-like filters resulting from training, similar to receptive fields in the primary visual cortex [31], are shown in Fig. 3 for both maximum likelihood and MPF solutions.

In summary, we have presented a novel, general purpose framework, called minimum probability flow, for fitting probabilistic models to data that outperforms current techniques in both learning time and accuracy. Our method works for any parametric model without hidden state variables, over either continuous or discrete state spaces, and we avoid explicit calculation of the partition function by employing deterministic dynamics in place of the slow sampling required by many existing approaches. Because MPF provides a simple and well-defined objective function, it can be minimized quickly using existing higher order gradient descent techniques. Furthermore, the objective function is convex for many models, including those in the exponential family, ensuring that the global minimum can be found with gradient descent. Finally, MPF is consistent—it will find the true parameter values when the data distribution belongs to the same family of parametric models as the model distribution.

We would like to thank Javier Movellan, Tamara Broderick, Miroslav Dudík, Gašper Tkačik, Robert E. Schapire, and William Bialek for sharing work in progress and data; Ashvin Vishwanath, Jonathon Shlens, Tony Bell, Charles Cadieu, Nicole Carlson, Christopher Hillar, Kilian Koepsell, Bruno Olshausen and the rest of the Redwood Center for many useful discussions; and the James S. McDonnell Foundation (M. R. D., P. B. B., J. S. D.) and

the Canadian Institute for Advanced Research—Neural Computation and Perception Program (J. S. D.) for financial support.

-
- [1] E. Schneidman *et al.*, *Nature (London)* **440**, 1007 (2006).
 - [2] J. Shlens *et al.*, *J. Neurosci.* **26**, 8254 (2006).
 - [3] A. Tang *et al.*, *J. Neurosci.* **28**, 505 (2008).
 - [4] O. Marre *et al.*, *Phys. Rev. Lett.* **102**, 138101 (2009).
 - [5] S. Yu *et al.*, *Cereb. Cortex* **18**, 2891 (2008).
 - [6] T. Broderick *et al.*, [arXiv:0712.2437](https://arxiv.org/abs/0712.2437).
 - [7] D. MacKay, *Information Theory, Inference and Learning Algorithms* (Cambridge University Press, Cambridge, England, 2002).
 - [8] I. C. Chou and E. O. Voit, *Math. Biosci.* **219**, 57 (2009).
 - [9] R. C. Aster, B. Borchers, and C. H. Thurber, *Parameter Estimation And Inverse Problems* (Elsevier Academic Press, Burlington, MA, 2005).
 - [10] S. Haykin, *Neural Networks And Learning Machines* (Prentice Hall, Englewood Cliffs, NJ, 2008), 3rd ed..
 - [11] A. J. Bell and T. J. Sejnowski, *Neural Comput.* **7**, 1129 (1995).
 - [12] T. Cover, J. Thomas, and J. Wiley, *Elements of Information Theory* (Wiley Online Library, New York, 1991), Vol. 1.
 - [13] R. Pathria, *Statistical Mechanics* (Pergamon Press, New York, 1972).
 - [14] T. Tanaka, *Phys. Rev. E* **58**, 2302 (1998).
 - [15] J. Fischer and K. H. Hertz, *Spin Glasses* (Cambridge University Press, Cambridge, England, 1991).
 - [16] H. Attias, *Adv. Neural Inf. Process. Syst.* **12** 209 (2000).
 - [17] J. M. Jaakkola, *Stat. Comput.*, **10**, 25 (2000).
 - [18] J. Besag, *Statistician*, **24**, 179 (1975).
 - [19] M. A. Carreira-Perpiñán and G. E. Hinton, University of Toronto, 2004 (unpublished).
 - [20] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *Cogn. Sci.* **9**, 147 (1985).
 - [21] A. Hyvärinen, *J. Mach. Learn. Res.* **6**, 695 (2005), <http://jmlr.csail.mit.edu/papers/volume6/hyvarinen05a/hyvarinen05a.pdf>.
 - [22] S. Lyu, in *Uncertainty in artificial intelligence: proceedings of the Twenty-fifth Conference (2009), June 18-21, 2009, Montreal, Quebec*, edited by J. Bilmes, A. Y. Ng, and D. McAllester (AUAI Press, Corvallis, OR, 2009).
 - [23] J. R. Movellan (unpublished).
 - [24] R. M. Neal, *Handbook of Markov Chain Monte Carlo* (Jan 2010), sections 5.2 and 5.3 for langevin dynamics.
 - [25] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.107.220601> for detailed derivations and additional information.
 - [26] J. Macke and S. Gerwin, Personal communication (2009).
 - [27] A. Hyvärinen, *Computational statistics & data analysis* **51**, 2499 (2007), .
 - [28] R. H. Swendsen and J. Wang, *Phys. Rev. Lett.* **58**, 86 (1987).
 - [29] <https://github.com/Sohl-Dickstein/Minimum-Probability-Flow-Learning>.
 - [30] D. J. Thouless, P. W. Anderson, R. G. Palmer, *Philos. Mag.* **35**, 593 (1977).
 - [31] P. Dayan and L. Abbott, *Theoretical Neuroscience* (MIT, Cambridge, MA, 2001), Vol. 83.

A new method for parameter estimation in probabilistic models: Minimum probability flow

Supplemental material

Appendix A: Taylor Expansion of KL Divergence

The minimum probability flow learning objective function $K(\theta)$ is found by taking up to the first order terms in the Taylor expansion of the KL divergence between the data distribution and the distribution resulting from running the dynamics for a time ϵ :

$$K(\theta) \approx D_{KL}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(t)}(\theta)) \Big|_{t=0} + \epsilon \frac{\partial D_{KL}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(t)}(\theta))}{\partial t} \Big|_{t=0} \quad (\text{A-1})$$

$$= 0 + \epsilon \frac{\partial D_{KL}(\mathbf{p}^{(0)} \parallel \mathbf{p}^{(t)}(\theta))}{\partial t} \Big|_{t=0} \quad (\text{A-2})$$

$$= \epsilon \frac{\partial}{\partial t} \left(\sum_{i \in \mathcal{D}} p_i^{(0)} \log \frac{p_i^{(0)}}{p_i^{(t)}} \right) \Big|_0 \quad (\text{A-3})$$

$$= -\epsilon \sum_{i \in \mathcal{D}} \frac{p_i^{(0)}}{p_i^{(0)}} \frac{\partial p_i^{(t)}}{\partial t} \Big|_0 \quad (\text{A-4})$$

$$= -\epsilon \sum_{i \in \mathcal{D}} \frac{\partial p_i^{(t)}}{\partial t} \Big|_0 \quad (\text{A-5})$$

$$= -\epsilon \left(\frac{\partial}{\partial t} \sum_{i \in \mathcal{D}} p_i^{(t)} \right) \Big|_0 \quad (\text{A-6})$$

$$= -\epsilon \frac{\partial}{\partial t} \left(1 - \sum_{i \notin \mathcal{D}} p_i^{(t)} \right) \Big|_0 \quad (\text{A-7})$$

$$= \epsilon \sum_{i \notin \mathcal{D}} \frac{\partial p_i^{(t)}}{\partial t} \Big|_0 \quad (\text{A-8})$$

$$= \epsilon \sum_{i \notin \mathcal{D}} \sum_{j \in \mathcal{D}} \Gamma_{ij} p_j^{(0)} \quad (\text{A-9})$$

$$= \frac{\epsilon}{|\mathcal{D}|} \sum_{i \notin \mathcal{D}} \sum_{j \in \mathcal{D}} \Gamma_{ij}, \quad (\text{A-10})$$

where we used the fact that $\sum_{i \in \mathcal{D}} p_i^{(t)} + \sum_{i \notin \mathcal{D}} p_i^{(t)} = 1$. This implies that the rate of growth of the KL divergence at time $t = 0$ equals the total initial flow of probability from states with data into states without.

Appendix B: Convexity

As observed by Macke and Gerwin [1], the MPF objective function is convex for models in the exponential family.

We wish to minimize

$$K = \sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{D}^c} \Gamma_{ij} p_i^{(0)}. \quad (\text{B-1})$$

K has derivative

$$\frac{\partial K}{\partial \theta_m} = \sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{D}^c} \left(\frac{\partial \Gamma_{ij}}{\partial \theta_m} \right) p_i^{(0)} \quad (\text{B-2})$$

$$= \frac{1}{2} \sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{D}^c} \Gamma_{ij} \left(\frac{\partial E_j}{\partial \theta_m} - \frac{\partial E_i}{\partial \theta_m} \right) p_i^{(0)}, \quad (\text{B-3})$$

and Hessian

$$\begin{aligned} \frac{\partial^2 K}{\partial \theta_m \partial \theta_n} &= \frac{1}{4} \sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{D}^c} \Gamma_{ij} \left(\frac{\partial E_j}{\partial \theta_m} - \frac{\partial E_i}{\partial \theta_m} \right) \left(\frac{\partial E_j}{\partial \theta_n} - \frac{\partial E_i}{\partial \theta_n} \right) p_i^{(0)} \\ &\quad + \frac{1}{2} \sum_{i \in \mathcal{D}} \sum_{j \in \mathcal{D}^c} \Gamma_{ij} \left(\frac{\partial^2 E_j}{\partial \theta_m \partial \theta_n} - \frac{\partial^2 E_i}{\partial \theta_m \partial \theta_n} \right) p_i^{(0)}. \end{aligned} \quad (\text{B-4})$$

The first term in the Hessian is a weighted sum of outer products, with non-negative weights $\frac{1}{4} \Gamma_{ij} p_i^{(0)}$, and is thus positive semidefinite. The second term is 0 for models in the exponential family (those with energy functions linear in their parameters).

Parameter estimation for models in the exponential family is therefore convex using minimum probability flow learning.

Appendix C: Relationship of MPF to other techniques

1. Score matching

Score matching, developed by Aapo Hyvärinen [2], is a method that learns parameters in a probabilistic model using only derivatives of the energy function evaluated over the data distribution (see Equation (C-5)). This sidesteps the need to explicitly sample or integrate over the model distribution. In score matching one minimizes the expected square distance of the score function with respect to spatial coordinates given by the data distribution from the similar score function given by the model distribution. A number of connections have been made between score matching and other learning techniques [3–6]. Here we show that in the correct limit, MPF also reduces to score matching.

For a d -dimensional, continuous state space, we can

write the MPF objective function as

$$\begin{aligned} K_{\text{MPF}} &= \frac{1}{N} \sum_{x \in \mathcal{D}} \int d^d y \Gamma(y, x) \\ &= \frac{1}{N} \sum_{x \in \mathcal{D}} \int d^d y g(y, x) e^{\frac{1}{2}(E(x|\theta) - E(y|\theta))}, \quad (\text{C-1}) \end{aligned}$$

where the sum $\sum_{x \in \mathcal{D}}$ is over all data samples, and N is the number of samples in the data set \mathcal{D} . Now we assume that transitions are only allowed from states x to states y that are within a hypercube of side length ϵ centered around x in state space. (The master equation will reduce to Gaussian diffusion as $\epsilon \rightarrow 0$.) Thus, the function $g(y, x)$ will equal 1 when y is within the x -centered cube (or x within the y -centered cube) and 0 otherwise. Calling this cube C_ϵ , and writing $y = x + \alpha$ with $\alpha \in C_\epsilon$, we have

$$K_{\text{MPF}} = \frac{1}{N} \sum_{x \in \mathcal{D}} \int_{C_\epsilon} d^d \alpha e^{\frac{1}{2}(E(x|\theta) - E(x+\alpha|\theta))}. \quad (\text{C-2})$$

If we Taylor expand in α to second order and ignore cubic and higher terms, we get

$$\begin{aligned} K_{\text{MPF}} &\approx \frac{1}{N} \sum_{x \in \mathcal{D}} \int_{C_\epsilon} d^d \alpha (1) \\ &\quad - \frac{1}{N} \sum_{x \in \mathcal{D}} \int_{C_\epsilon} d^d \alpha \frac{1}{2} \sum_{i=1}^d \alpha_i \nabla_{x_i} E(x|\theta) \\ &\quad + \frac{1}{N} \sum_{x \in \mathcal{D}} \int_{C_\epsilon} d^d \alpha \frac{1}{4} \left(\frac{1}{2} \left[\sum_{i=1}^d \alpha_i \nabla_{x_i} E(x|\theta) \right]^2 \right. \\ &\quad \left. - \sum_{i,j=1}^d \alpha_i \alpha_j \nabla_{x_i} \nabla_{x_j} E(x|\theta) \right). \quad (\text{C-3}) \end{aligned}$$

This reduces to

$$\begin{aligned} K_{\text{MPF}} &\approx \frac{1}{N} \sum_{x \in \mathcal{D}} \left[\epsilon^d + \frac{1}{4} \left(\frac{1}{2} \frac{1}{12} \epsilon^{d+2} \sum_{i=1}^d \left[\nabla_{x_i} E(x|\theta) \right]^2 \right. \right. \\ &\quad \left. \left. - \frac{1}{12} \epsilon^{d+2} \sum_{i=1}^d \nabla_{x_i}^2 E(x|\theta) \right) \right], \quad (\text{C-4}) \end{aligned}$$

which, removing a constant offset and scaling factor, is exactly equal to the score matching objective function,

$$K_{\text{MPF}} \sim \frac{1}{N} \sum_{x \in \mathcal{D}} \left[\frac{1}{2} \nabla E(x|\theta) \cdot \nabla E(x|\theta) - \nabla^2 E(x|\theta) \right] \quad (\text{C-5})$$

$$= K_{\text{SM}}. \quad (\text{C-6})$$

Score matching is thus equivalent to MPF when the connectivity function $g(y, x)$ is non-zero only for states infinitesimally close to each other. It should be noted that the score matching estimator has a closed-form solution when the model distribution belongs to the exponential family [7], so the same can be said for MPF in this limit.

2. Contrastive divergence

Contrastive divergence [8, 9] is a variation on steepest gradient descent of the maximum (log) likelihood (ML) objective function. Rather than integrating over the full model distribution, CD approximates the partition function term in the gradient by averaging over the distribution obtained after taking a few, or only one, Markov chain Monte Carlo (MCMC) step away from the data distribution (Equation C-7). Qualitatively, one can imagine that the data distribution is contrasted against a distribution which has evolved only a small distance towards the model distribution, whereas it would be contrasted against the true model distribution in traditional MCMC approaches. Although CD is not guaranteed to converge to the right answer, or even to a fixed point, it has proven to be an effective and fast heuristic for parameter estimation [10, 11].

The contrastive divergence update rule can be written in the form

$$\Delta \theta_{CD} \propto - \sum_{j \in \mathcal{D}} \sum_{i \notin \mathcal{D}} \left[\frac{\partial E_j(\theta)}{\partial \theta} - \frac{\partial E_i(\theta)}{\partial \theta} \right] T_{ij}, \quad (\text{C-7})$$

where T_{ij} is the probability of transitioning from state j to state i in a single Markov chain Monte Carlo step (or a small number of steps). Equation C-7 has obvious similarities to the MPF learning gradient

$$\begin{aligned} \frac{\partial K(\theta)}{\partial \theta} &= \frac{\epsilon}{2N} \sum_{j \in \mathcal{D}} \sum_{i \notin \mathcal{D}} \left[\frac{\partial E_j(\theta)}{\partial \theta} - \frac{\partial E_i(\theta)}{\partial \theta} \right] \\ &\quad g_{ij} \exp \left[\frac{1}{2} (E_j(\theta) - E_i(\theta)) \right]. \quad (\text{C-8}) \end{aligned}$$

Thus, steepest gradient descent under MPF resembles CD updates, but with the MCMC sampling/rejection step T_{ij} replaced by a weighting factor $g_{ij} \exp \left[\frac{1}{2} (E_j(\theta) - E_i(\theta)) \right]$. Note that this difference in form provides MPF with a well-defined objective function, and it guarantees consistency (*i.e.*, there is a global minimum when model and data distributions agree).

Appendix D: Sampling the connectivity matrix Γ

The MPF learning scheme is blind to regions in state space which data states don't have any connectivity to - the flow at time 0 is only a function of the states that are directly connected to data states. To get the most informative learning signal, it seems sensible to encourage probability flow directly between data states and states that are probable under the model. That way the objective function is sensitive to the regions which are probable under the model. We believe nearest neighbor connectivity schemes are effective largely because as the parameters converge the regions around data states become the high probability regions for the model. We wish to try connectivity schemes other than nearest neighbors to allow probability to most efficiently flow between data states and high probability model states. In order to do

so we need to slightly extend the MPF algorithm. We do this by allowing the connectivity pattern in Γ to be sampled independently in every infinitesimal time step.

Since Γ is now sampled, we will modify detailed balance to demand that, averaging over the choices for Γ , the net flow between pairs of states is 0.

$$\langle \Gamma_{ji} p_i^{(\infty)}(\theta) \rangle = \langle \Gamma_{ij} p_j^{(\infty)}(\theta) \rangle \quad (\text{D-1})$$

$$\langle \Gamma_{ji} \rangle p_i^{(\infty)}(\theta) = \langle \Gamma_{ij} \rangle p_j^{(\infty)}(\theta), \quad (\text{D-2})$$

where the ensemble average is over the connectivity scheme for Γ . We describe the connectivity scheme via a function g_{ij} , such that the probability of there being a connection from state j to state i at any given moment is g_{ij} . We also introduce a function F_{ij} , which provides the value Γ_{ij} takes on when a connection occurs from j to i . That is, it is the probability flow rate when flow occurs -

$$\langle \Gamma_{ij} \rangle = g_{ij} F_{ij}. \quad (\text{D-3})$$

Detailed balance now becomes

$$g_{ji} F_{ji} p_i^{(\infty)}(\theta) = g_{ij} F_{ij} p_j^{(\infty)}(\theta). \quad (\text{D-4})$$

Solving for \mathbf{F} we find

$$\frac{F_{ij}}{F_{ji}} = \frac{g_{ji} p_i^{(\infty)}(\theta)}{g_{ij} p_j^{(\infty)}(\theta)} = \frac{g_{ji}}{g_{ij}} \exp[E_j(\theta) - E_i(\theta)]. \quad (\text{D-5})$$

\mathbf{F} is underconstrained by the above equation. Motivated by symmetry and aesthetics, we choose as the form for the (non-zero, non-diagonal) entries in \mathbf{F}

$$F_{ij} = \left(\frac{g_{ji}}{g_{ij}} \right)^{\frac{1}{2}} \exp \left[\frac{1}{2} (E_j(\theta) - E_i(\theta)) \right]. \quad (\text{D-6})$$

Γ is now populated as

$$r_{ij} \sim \text{rand}[0, 1] \quad (\text{D-7})$$

$$\Gamma_{ij} = \begin{cases} -\sum_{k \neq i} \Gamma_{ki} & i = j \\ F_{ij} & r_{ij} < g_{ij} \text{ and } i \neq j \\ 0 & r_{ij} \geq g_{ij} \text{ and } i \neq j \end{cases}. \quad (\text{D-8})$$

Similarly, its average value can be written as

$$\langle \Gamma_{ij} \rangle = g_{ij} \left(\frac{g_{ji}}{g_{ij}} \right)^{\frac{1}{2}} \exp \left[\frac{1}{2} (E_j(\theta) - E_i(\theta)) \right] \quad (\text{D-9})$$

$$= (g_{ij} g_{ji})^{\frac{1}{2}} \exp \left[\frac{1}{2} (E_j(\theta) - E_i(\theta)) \right]. \quad (\text{D-10})$$

So, we can use any connectivity scheme \mathbf{g} in learning. We just need to scale the non-zero, non-diagonal entries in Γ by $\left(\frac{g_{ji}}{g_{ij}} \right)^{\frac{1}{2}}$ so as to compensate for the biases introduced by the connectivity scheme.

The full MPF objective function in this case is

$$K = \sum_{j \in \mathcal{D}} \sum_{i \notin \mathcal{D}} g_{ij} \left(\frac{g_{ji}}{g_{ij}} \right)^{\frac{1}{2}} \exp \left[\frac{1}{2} (E_j - E_i) \right] \quad (\text{D-11})$$

where the inner sum is found by averaging over samples from g_{ij} .

TABLE I. Mean square error in recovered coupling strengths (ϵ_J), mean square error in pairwise correlations (ϵ_{corr}) and learning time for MPF versus mean field theory with TAP correction (MFT+TAP), 1-step and 10-step contrastive divergence (CD-1 and CD-10), and pseudolikelihood (PL).

TECHNIQUE	ϵ_J	ϵ_{corr}	TIME (S)
MPF	0.0172	0.0025	~60
MFT+TAP	7.7704	0.0983	0.1
CD-1	0.3196	0.0127	~20000
CD-10	0.3341	0.0123	~20000
PL	0.0582	0.0036	~800

Appendix E: Additional information on Ising spin glass example from main text

1. Competing techniques

The four competing techniques against which MPF was compared are: mean field theory (MFT) with Thouless-Anderson-Palmer (TAP) corrections [12], one-step and ten-step contrastive divergence [9] (CD-1 and CD-10), and pseudolikelihood [13]. Table I shows the relative performance at convergence for each technique in terms of convergence time and mean square error in coupling strengths and pairwise correlations.

MFT with TAP involves approximating the Gibbs free energy of the model with a second-order Plefka expansion [14]. MFT+TAP is fast because it involves only an inversion of the magnetic susceptibility matrix, but it can perform poorly, for instance near criticality [15].

Contrastive divergence approximates the term involving the partition function in $\partial_{\theta} D_{KL}(\mathbf{p}^{(0)} \| \mathbf{p}^{(\infty)}(\theta))$, via a Markov chain which is initialized at the data distribution $\mathbf{p}^{(0)}$, and then truncated after only a small number of sampling steps. It is commonly used in machine learning, and provides an effective and fast stochastic parameter update rule for learning in many probabilistic models. However, it is not guaranteed to converge to a fixed point, and it does not correspond exactly to an objective function. The relationship between our technique and contrastive divergence is discussed in the Supplemental Material.

Pseudolikelihood approximates the joint probability distribution of a collection of random variables with a product of conditional distributions, where each factor is the distribution of a single random variable conditioned on the others:

$$p(x_1, x_2, \dots, x_d) \rightarrow \prod_{i=1}^d p(x_i | x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_d) \quad (\text{E-1})$$

This approach often leads to surprisingly good estimates, despite the extreme nature of the approximation.

2. Optimization steps taken for parameter estimation algorithms

a. Minimum Probability Flow and Pseudolikelihood

Both Minimum Probability Flow and Pseudolikelihood have well defined objective functions and gradients. Parameter estimation was thus performed by applying an off the shelf L-BFGS (quasi-Newton gradient descent) implementation [16] to their objective functions evaluated over the full training dataset \mathcal{D} .

3. Contrastive Divergence

The CD update rule was computed using the full training dataset. The learning rate was annealed in a linear fashion from 3.0 to 0.1 to accelerate convergence.

a. Mean Field Theory

Mean field theory requires the computation of the inverse of the magnetic susceptibility matrix, which, for strong correlations, was often singular. A regularized pseudoinverse was used in the following manner:

$$A = (\chi^T \chi + \lambda I)^+ \chi^T, \quad (\text{E-2})$$

where I is the identity matrix, M^+ denotes the Moore-Penrose pseudoinverse of a matrix M , χ is the magnetic susceptibility $\chi_{ij} = \langle x_i x_j \rangle - \langle x_i \rangle \langle x_j \rangle$, and λ is a regularizing parameter. This technique is known as stochastic robust approximation [17].

4. Dependence of computation time on sample size

For the Ising spin glass example described above and in the text, we measured both the time to evaluate the objective function and the time for the L-BFGS implementation in Section E2a to converge as a function of batch size. As can be seen in Figure E-1, for large batch size the objective function evaluation time is linear, and the convergence time is approximately linear.

Appendix F: Additional Ising spin glass comparison

The Ising model has a long and storied history in physics [18] and machine learning [19] and it has recently been found to be a surprisingly useful model for networks of neurons in the retina [20, 21]. The ability to fit Ising models to the activity of large groups of simultaneously recorded neurons is of current interest given the increasing number of these types of data sets from the retina, cortex and other brain structures.

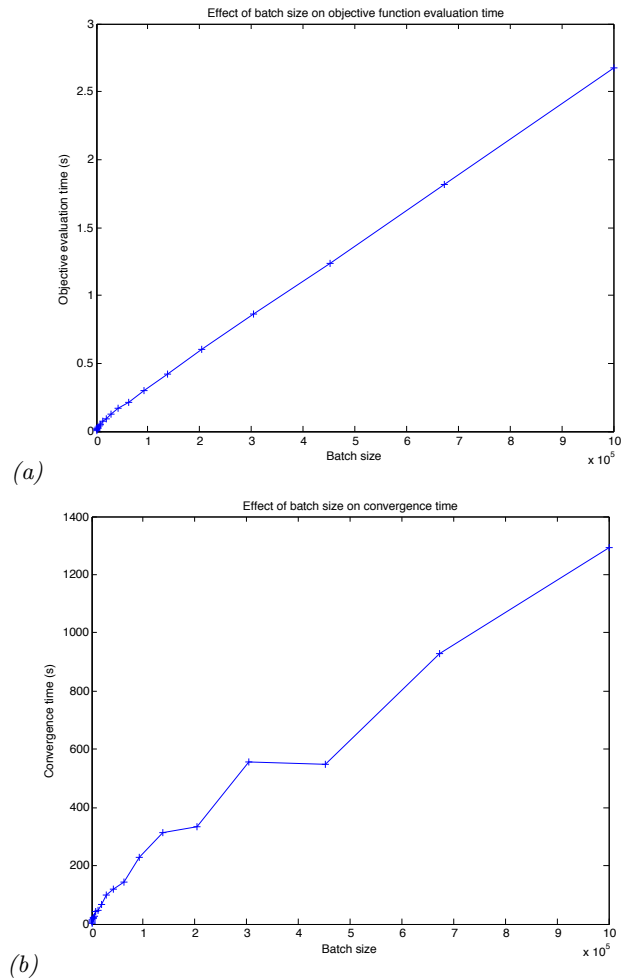


FIG. E-1. The time taken for (a) evaluation of the MPF objective function and (b) convergence of the L-BFGS parameter estimation algorithm as a function of training batch size for the Ising spin glass model presented in the text. Parameter estimation involves many parameter update steps, each of which requires reevaluating the MPF objective function and gradient.

We fit an Ising model (fully visible Boltzmann machine) of the form

$$p^{(\infty)}(\mathbf{x}; \mathbf{J}) = \frac{1}{Z(\mathbf{J})} \exp \left[- \sum_{i,j} J_{ij} x_i x_j \right] \quad (\text{F-1})$$

to a set of N d -element iid data samples $\{x^{(i)} | i = 1 \dots N\}$ generated via Gibbs sampling from an Ising model as described below, where each of the d elements of \mathbf{x} is either 0 or 1. Because each $x_i \in \{0, 1\}$, $x_i^2 = x_i$, we can write the energy function as

$$E(\mathbf{x}; \mathbf{J}) = \sum_{i,j \neq i} J_{ij} x_i x_j + \sum_i J_{ii} x_i. \quad (\text{F-2})$$

The probability flow matrix $\mathbf{\Gamma}$ has $2^N \times 2^N$ elements,




FIG. F-1. A demonstration of rapid fitting of the Ising model by minimum probability flow learning. The mean absolute error in the learned model’s correlation matrix is shown as a function of learning time for a 40 unit fully connected Ising model. Convergence is reached in about 15 seconds for 20,000 samples.

but for learning we populate it extremely sparsely, setting

$$g_{ij} = g_{ji} = \begin{cases} 1 & \text{states } i \text{ and } j \text{ differ by single bit flip} \\ 0 & \text{otherwise} \end{cases}. \quad (\text{F-3})$$

Figure F-1 shows the average error in predicted correlations as a function of learning time for 20,000 samples from a 40 unit, fully connected Ising model. The final absolute correlation error is 0.0058. The J_{ij} used were graciously provided by Broderick and coauthors, and were identical to those used for synthetic data generation in the 2008 paper “Faster solutions of the inverse pairwise Ising problem” [22]. Training was performed on 20,000 samples so as to match the number of samples used in section III.A. of Broderick et al. Note that given sufficient samples, the minimum probability flow algorithm would converge exactly to the right answer, as learning in the Ising model is convex (see Appendix B), and has its global minimum at the true solution. On an 8 core 2.33 GHz Intel Xeon, the learning converges in about 15 seconds. Broderick et al. perform a similar learning task on a 100-CPU grid computing cluster, with a convergence time of approximately 200 seconds.

Appendix G: Continuous state space independent component analysis (ICA) [23] model

Training was performed on 100,000 10×10 pixel whitened natural image patches from the van Hateren database [24]. Minimization was performed by alternating between minimization of the objective function in Equation D-11 and updates to the continuous state space connectivity function $g(\mathbf{x}_j, \mathbf{x}_i)$, as described in more detail in Section H. Both training techniques were initial-

ized with identical isotropic Gaussian noise (with variance 0.01, such that each receptive field was initialized to nearly unit length), and trained on the same image patches, which accounts for the similarity of individual filters found by the algorithms.

Appendix H: Continuous state space learning with the connectivity function set via Hamiltonian Monte Carlo

Choosing the connectivity matrix g_{ij} for Minimum Probability Flow Learning is relatively straightforward in systems with binary or discrete state spaces. Nearly any nearest neighbor style scheme seems to work quite well. In continuous state spaces $\mathbf{q} \in \mathbb{R}^d$ however, connectivity functions $g(\mathbf{q}_i, \mathbf{q}_j)$ based on nearest neighbors prove insufficient. For instance, if the non-zero entries in $g(\mathbf{q}_i, \mathbf{q}_j)$ are drawn from an isotropic Gaussian centered on \mathbf{q}_j , then several hundred non-zero $g(\mathbf{q}_i, \mathbf{q}_j)$ are required for every value of \mathbf{q}_j in order to achieve effective parameter estimation in some fairly standard problems, such as receptive field estimation in Independent Component Analysis [23].

Qualitatively, we desire to connect every data state $\mathbf{q}_j \in \mathcal{D}$ to the non data states \mathbf{q}_i which will be most informative for learning. The most informative states are those which have high probability under the model distribution $p^{(\infty)}(\mathbf{q})$. We therefore propose to populate $g(\mathbf{q}_i, \mathbf{q}_j)$ using a Markov transition function for the model distribution. Borrowing techniques from Hamiltonian Monte Carlo [25] we use Hamiltonian dynamics in our transition function, so as to effectively explore the state space.

1. Extending the state space

In order to implement Hamiltonian dynamics, we first extend the state space to include auxiliary momentum variables.

The initial data and model distributions are $p^{(0)}(\mathbf{q})$ and

$$p^{(\infty)}(\mathbf{q}; \theta) = \frac{\exp(-E(\mathbf{q}; \theta))}{Z(\theta)}. \quad (\text{H-1})$$

with state space $\mathbf{q} \in \mathbb{R}^d$. We introduce auxiliary momentum variables $\mathbf{v} \in \mathbb{R}^d$ for each state variable \mathbf{q} , and call the extended state space including the momentum variables $\mathbf{x} = \{\mathbf{q}, \mathbf{v}\}$. The momentum variables are given an isotropic gaussian distribution,

$$p(\mathbf{v}) = \frac{\exp(-\frac{1}{2}\mathbf{v}^T\mathbf{v})}{\sqrt{2\pi}}, \quad (\text{H-2})$$

and the extended data and model distributions become

$$p^{(0)}(\mathbf{x}) = p^{(0)}(\mathbf{q})p(\mathbf{v}) \quad (\text{H-3})$$

$$= p^{(0)}(\mathbf{q}) \frac{\exp\left(-\frac{1}{2}\mathbf{v}^T\mathbf{v}\right)}{\sqrt{2\pi}} \quad (\text{H-4})$$

$$p^{(\infty)}(\mathbf{x}; \theta) = p^{(\infty)}(\mathbf{q}; \theta)p(\mathbf{v}) \quad (\text{H-5})$$

$$= \frac{\exp(-E(\mathbf{q}; \theta)) \exp\left(-\frac{1}{2}\mathbf{v}^T\mathbf{v}\right)}{Z(\theta) \sqrt{2\pi}} \quad (\text{H-6})$$

$$= \frac{\exp(-H(\mathbf{x}; \theta))}{Z(\theta) \sqrt{2\pi}} \quad (\text{H-7})$$

$$H(\mathbf{x}; \theta) = E(\mathbf{q}; \theta) + \frac{1}{2}\mathbf{v}^T\mathbf{v}. \quad (\text{H-8})$$

The initial (data) distribution over the joint space \mathbf{x} can be realized by drawing a momentum \mathbf{v} from a uniform Gaussian distribution for every observation \mathbf{q} in the dataset \mathcal{D} .

2. Defining the connectivity function $g(\mathbf{x}_i, \mathbf{x}_j)$

We connect every state \mathbf{x}_j to all states which satisfy one of the following 2 criteria,

1. All states which share the same position \mathbf{q}_j , with a quadratic falloff in $g(\mathbf{x}_i, \mathbf{x}_j)$ with the momentum difference $\mathbf{v}_i - \mathbf{v}_j$.
2. The state which is reached by simulating Hamiltonian dynamics for a fixed time t on the system described by $H(\mathbf{x}; \theta_H)$, and then negating the momentum. Note that the parameter vector θ_H is used only for the Hamiltonian dynamics.

More formally,

$$g(\mathbf{x}_i, \mathbf{x}_j) = \delta(\mathbf{q}_i - \mathbf{q}_j) \exp\left(-\|\mathbf{v}_i - \mathbf{v}_j\|_2^2\right) + \delta(\mathbf{x}_i - \text{HAM}(\mathbf{x}_j; \theta_H)) \quad (\text{H-9})$$

where if $\mathbf{x}' = \text{HAM}(\mathbf{x}; \theta_H)$, then \mathbf{x}' is the state that results from integrating Hamiltonian dynamics for a time t and then negating the momentum. Because of the momentum negation, $\mathbf{x} = \text{HAM}(\mathbf{x}'; \theta_H)$, and $g(\mathbf{x}_i, \mathbf{x}_j) = g(\mathbf{x}_j, \mathbf{x}_i)$.

3. Discretizing Hamiltonian dynamics

It is generally impossible to **exactly** simulate the Hamiltonian dynamics for the system described by $H(\mathbf{x}; \theta_H)$. However, if $\text{HAM}(\mathbf{x}; \theta_H)$ is set to simulate Hamiltonian dynamics via a series of leapfrog steps, it retains the important properties of reversibility and phase space volume conservation, and can be used in the connectivity function $g(\mathbf{x}_i, \mathbf{x}_j)$ in Equation H-9. In practice, therefore, $\text{HAM}(\mathbf{x}; \theta_H)$ involves the simulation of Hamiltonian dynamics by a series of leapfrog steps.

4. MPF objective function

The MPF objective function for continuous state spaces and a list of observations \mathcal{D} is

$$K(\theta; \mathcal{D}, \theta_H) = \sum_{\mathbf{x}_j \in \mathcal{D}} \int g(\mathbf{x}_i, \mathbf{x}_j) \exp\left(\frac{1}{2}[H(\mathbf{x}_j; \theta) - H(\mathbf{x}_i; \theta)]\right) d\mathbf{x}_i. \quad (\text{H-10})$$

For the connectivity function $g(\mathbf{x}_i, \mathbf{x}_j)$ given in Section H2, this reduces to

$$K(\theta; \mathcal{D}, \theta_H) = \sum_{\mathbf{x}_j \in \mathcal{D}} \int \exp\left(-\|\mathbf{v}_i - \mathbf{v}_j\|_2^2\right) \exp\left(\frac{1}{2}\left[\frac{1}{2}\mathbf{v}_j^T\mathbf{v}_j - \frac{1}{2}\mathbf{v}_i^T\mathbf{v}_i\right]\right) d\mathbf{v}_i + \sum_{\mathbf{x}_j \in \mathcal{D}} \exp\left(\frac{1}{2}[H(\mathbf{x}_j; \theta) - H(\text{HAM}(\mathbf{x}_j; \theta_H); \theta)]\right). \quad (\text{H-11})$$

Note that the first term does not depend on the parameters θ , and is thus just a constant offset which can be ignored during optimization. Therefore, we can say

$$K(\theta; \mathcal{D}, \theta_H) \sim \sum_{\mathbf{x}_j \in \mathcal{D}} \exp\left(\frac{1}{2}[H(\mathbf{x}_j; \theta) - H(\text{HAM}(\mathbf{x}_j; \theta_H); \theta)]\right). \quad (\text{H-12})$$

Parameter estimation is performed by finding the parameter vector $\hat{\theta}$ which minimizes the objective function $K(\theta; \mathcal{D}, \theta_H)$,

$$\hat{\theta} = \underset{\theta}{\text{argmin}} K(\theta; \mathcal{D}, \theta_H). \quad (\text{H-13})$$

5. Iteratively improving the objective function

The more similar θ_H is to θ , the more informative $g(\mathbf{x}_i, \mathbf{x}_j)$ is for learning. If θ_H and θ are dissimilar, then many more data samples will be required in \mathcal{D} to effectively learn. Therefore, we iterate the following procedure, which alternates between finding the $\hat{\theta}$ which minimizes $K(\theta; \mathcal{D}, \theta_H)$, and improving θ_H by setting it to $\hat{\theta}$,

1. Set $\hat{\theta}^{t+1} = \underset{\theta}{\text{argmin}} K(\theta; \mathcal{D}, \theta_H^t)$
2. Set $\theta_H^{t+1} = \hat{\theta}^{t+1}$

$\hat{\theta}^t$ then represents a steadily improving estimate for the parameter values which best fit the model distribution

$p^{(\infty)}(\mathbf{q}; \theta)$ to the data distribution $p^{(0)}(\mathbf{q})$, described by observations \mathcal{D} . Practically, step 1 above will frequently

be truncated early, perhaps after 10 or 100 L-BFGS gradient descent steps.

-
- [1] J. Macke and S. Gerwinn, Personal communication(2009).
 - [2] A. Hyvärinen, Journal of Machine Learning Research **6**, 695 (2005).
 - [3] A. Hyvärinen, IEEE Transactions on Neural Networks(Jan 2007).
 - [4] J. Sohl-Dickstein and B. Olshausen, Redwood Center Technical Report(2009).
 - [5] J. R. Movellan, unpublished draft(Jan 2008).
 - [6] S. Lyu, The proceedings of the 25th conference on uncertainty in artificial intelligence (UAI*90)(2009).
 - [7] A. Hyvärinen, Computational statistics & data analysis **51**, 2499 (2007), ISSN 0167-9473.
 - [8] M. Welling and G. Hinton, Lecture Notes in Computer Science(Jan 2002).
 - [9] M. A. Carreira-Perpiñán and G. E. Hinton, Technical report, Dept. of Computer Science, University of Toronto(2004).
 - [10] D. MacKay, Failures of the one-step learning algorithm(Jan 2001).
 - [11] A. Yuille, Department of Statistics, UCLA. Department of Statistics Papers.(2005).
 - [12] P. R. Thouless DJ, Anderson PW, Philos. Mag. 35 p.593(1977).
 - [13] J. Besag, The Statistician, 24(3), 179-195(1975).
 - [14] P. T, J. Phys. A: Math. Gen. 15 1971(1982).
 - [15] J. Fischer KH, Hertz, *Spin Glasses* (Cambridge University Press, 1991).
 - [16] M. Schmidt, <http://www.cs.ubc.ca/~schmidtm/Software/minFunc.html> (2005).
 - [17] S. Boyd and L. Vandenberghe, *Convex optimization* (Cambridge Univ Pr, 2004) ISBN 0521833787.
 - [18] S. G. Brush, Reviews of Modern Physics **39**, 883 (Oct 1967).
 - [19] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, Cognitive Science **9**, 147 (1985).
 - [20] E. Schneidman, M. J. B. 2nd, R. Segev, and W. Bialek, Nature **440**, 1007 (2006).
 - [21] J. Shlens, G. D. Field, J. L. Gauthier, M. I. Grivich, D. Petrusca, A. Sher, A. M. Litke, and E. J. Chichilnisky, J. Neurosci. **26**, 8254 (2006).
 - [22] T. Broderick, M. Dudík, G. Tkačik, R. Schapire, and W. Bialek, E-print arXiv(Jan 2007).
 - [23] S. T. Bell AJ, Neural Computation 1995; vol. 7:1129-1159(1995).
 - [24] J. H. v. Hateren and A. v. d. Schaaf, Proceedings: Biological Sciences **265**, 359 (Mar 1998).
 - [25] R. M. Neal, Handbook of Markov Chain Monte Carlo(Jan 2010), sections 5.2 and 5.3 for langevin dynamics.