Bayesian approach to inverse statistical mechanics

Michael Habeck*

Institute for Mathematical Stochastics, University of Göttingen, Goldschmidtstrasse 7, 37077 Göttingen, Germany

(Received 27 December 2013; published 9 May 2014)

Inverse statistical mechanics aims to determine particle interactions from ensemble properties. This article looks at this inverse problem from a Bayesian perspective and discusses several statistical estimators to solve it. In addition, a sequential Monte Carlo algorithm is proposed that draws the interaction parameters from their posterior probability distribution. The posterior probability involves an intractable partition function that is estimated along with the interactions. The method is illustrated for inverse problems of varying complexity, including the estimation of a temperature, the inverse Ising problem, maximum entropy fitting, and the reconstruction of molecular interaction potentials.

DOI: 10.1103/PhysRevE.89.052113 PACS number(s): 02.50.Tt, 05.10.−a, 02.70.Rr

I. INTRODUCTION

Standard applications of statistical mechanics predict ensemble properties from microscopic interactions. Inverse problems of statistical mechanics involve "the determination of the interactions between the particles from the properties of the system" [1]. Problems of inverse statistical mechanics are important in diverse application areas. Inverse Ising and Potts methods abound in neuroscience and biology [2–5]. In materials science inverse statistical mechanics is used to design self-assembling materials with a predefined ground state [6,7]. Another instance is the estimation of coarse-grained models for biological macromolecules [8–10].

The simplest inverse problem in statistical mechanics involves the estimation of a single parameter, the system's temperature in the canonical ensemble. Mandelbrot and others have discussed the temperature-estimation problem from the point of view of statistical inference [11–13]. Specific instances of inverse statistical mechanics problems have been given particular attention. The inverse Ising problem, for example, has been addressed by a diverse array of techniques including naive mean-field theory (nMF) [14], pseudolikelihood maximization (PLM) [15], message passing [16], and minimum probability flow learning [17].

This article generalizes previous statistical analyses of inverse problems in statistical mechanics and adds a Bayesian perspective. A physical interpretation of the statistical estimators is discussed. A sequential Monte Carlo sampler is developed to apply the formalism in practice and demonstrated on various challenging inverse problems including the estimation of the temperature of a canonical ensemble, the inverse Ising problem, and the estimation of an interaction potential. Moreover the relation to the maximum entropy method [18] is highlighted and demonstrated.

II. ESTIMATION OF INTERACTION PARAMETERS

In the following the Hamiltonian of the system is assumed to be

$$E(x) = \lambda \cdot f(x) = \sum_{k=1}^{K} \lambda_k f_k(x) \tag{1}$$

where x is the system's configuration, $f_k(x)$ are K features or collective variables that encode the configuration, and λ_k are the associated interaction parameters. The Hamiltonian is assumed to be linear in the interaction parameters. This family of Hamiltonians covers a broad range of physical systems and interactions. However, also nonlinear extensions could be treated with the Bayesian formalism outlined below but are beyond the scope of this article.

Examples of physical systems that follow a linear Hamiltonian include the Ising model. This system has only a single feature $f_1(x) = -\sum_{\langle ij \rangle} x_i x_j$ with associated parameter $\lambda_1 = J$, where x is a vector of spin orientations and J is the coupling between neighboring spins on the two-dimensional lattice. If an external field of strength h is applied, the parameters are $\lambda = (J,h)$, and the magnetization $f_2(x) = -\sum_i x_i$ becomes a second collective variable. In case of a fluid governed by Lennard-Jones (LJ) interactions

$$E_{\mathrm{LJ}}(r) = 4\epsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right],$$

the features and interaction parameters are $f_k(x) = (-)^k \sum_{i<j} \|x_i - x_j\|^{-6k}$ and $\lambda_k = 4\beta\epsilon\sigma^{6k}$, $k = 1,2$, where β is the inverse temperature and ϵ and σ are the standard LJ parameters.

Configurations follow the Gibbs ensemble

$$p(x|\lambda) = \frac{\rho(x)}{Z(\lambda)}\, e^{-\lambda \cdot f(x)}, \tag{2}$$

where $\rho(x)$ is a reference distribution, for example, a uniform distribution over a finite box. The reference distribution also allows us to incorporate knowledge about an unperturbed Hamiltonian E_0 in which case $\rho(x) = \exp\{-E_0(x)\}$ and to estimate a perturbation $E(x)$. $Z(\lambda)$ is the partition function

$$Z(\lambda) = \int \rho(x)\, e^{-\lambda \cdot f(x)}\, dx \tag{3}$$

and generates the moments of the features: $-\nabla_\lambda \ln Z(\lambda) = \langle f \rangle_\lambda$, where $\langle \cdot \rangle_\lambda$ indicates an average over the Gibbs ensemble (2). For complex systems the partition function is typically unknown and cannot be evaluated by simple numerical means such as quadrature.

*mhabeck@gwdg.de

A. Statistical estimators and physical analogs

The inverse problem that we aim to solve is to recover the interaction parameters from observations of the system. The first step is to determine the likelihood of observing features f given the interaction parameters λ:

$$p(f|\lambda) = \int p(f|x)p(x|\lambda)dx = \frac{g(f)}{Z(\lambda)}\, e^{-\lambda\cdot f}, \tag{4}$$

where we introduce the distribution

$$g(f) = \int \left[\prod_k \delta(f_k - f_k(x))\right] \rho(x)\, dx, \tag{5}$$

with δ denoting the Dirac δ function. $g(f)$ is the projection of the reference distribution $\rho(x)$ into feature space. It measures the degeneracy of the collective variables and can be viewed as a multidimensional density of states. With the help of the feature distribution we can write the partition function (3) as an integral over the collective variables:

$$Z(\lambda) = \int g(f)\, e^{-\lambda\cdot f}\, df. \tag{6}$$

Probability density (4) is the sampling distribution of the features and a natural member of the exponential family [11]. Given the interaction parameters λ, the most likely features are determined by maximizing $p(f|\lambda)$ as a function of f. The resulting equation

$$\nabla_f \ln g(f) = \lambda \tag{7}$$

relates the interaction parameters directly to the collective variables. An alternative is to search for the Hamiltonian under which the observed features are highly probable. This is obtained by maximizing $p(f|\lambda)$ as a function of the interactions leading to

$$-\nabla_\lambda \ln Z(\lambda) = f. \tag{8}$$

Equations (7) and (8) have the same structure, only the role of the conjugate variables f and λ is interchanged.

If $\lambda = \beta$ is the inverse temperature and $f(x) = E(x)$ is the Hamiltonian itself, the above equations correspond to the microcanonical and canonical definition of the temperature. Equation (7) defines the microcanonical temperature β as a function of the energy E, whereas Eq. (8) defines the temperature implicitly in terms of the internal energy $\langle E\rangle_\beta = -\partial_\beta \ln Z(\beta)$. Interestingly expression (7) is an unbiased estimator of the interactions [13]:

$$\begin{aligned}\langle \nabla_f \ln g(f)\rangle_\lambda &= \int [\nabla_f g(f)]\, \frac{e^{-\lambda\cdot f}}{Z(\lambda)}\, df \\ &= -\frac{1}{Z(\lambda)} \int g(f)\, [\nabla_f\, e^{-\lambda\cdot f}]\, df = \lambda,\end{aligned} \tag{9}$$

where partial integration gives the second last equality.

We now assume that N realizations of the features f_n have been measured, for example, by observing the system's configuration N times in which case $f_n = f(x_n)$. [Note that f_n is a vector with elements $f_{nk} = f_k(x_n)$ and is different from f_k, which is a scalar function.] One way to determine the interaction parameters is to evaluate the microcanonical estimator (7) at the sample average of the features $\overline{f} = \frac{1}{N}\sum_{n=1}^N f_n$:

$$\hat{\lambda}_{\text{micro}} = \nabla_f \ln g(f)\,\big|_{f=\overline{f}}\,. \tag{10}$$

Unfortunately we can only rarely use this estimator in practice because $g(f)$ is unknown for most systems. In case of systems with a continuous configuration space it is possible to approximate the microcanonical temperature by the configurational temperature [19,20]. Using the configurational temperature formalism one can then derive a system of linear equations for the interaction parameters [10].

The maximum likelihood (ML) approach [11] maximizes the probability of observing all N feature vectors as a function of the interaction parameters λ:

$$p(f_1,\ldots,f_N|\lambda) = \prod_{n=1}^N p(f_n|\lambda) \propto \left(\frac{e^{-\lambda\cdot\overline{f}}}{Z(\lambda)}\right)^N =: L(\lambda). \tag{11}$$

The likelihood function $L(\lambda)$ depends on the observed features only through the sufficient statistic $\overline{f}$. Maximization of the likelihood function gives an implicit equation for the ML estimator $\hat{\lambda}_{\text{ML}}$:

$$\overline{f} = -\nabla_\lambda \ln Z(\lambda)|_{\lambda=\hat{\lambda}_{\text{ML}}} = \langle f\rangle_{\lambda=\hat{\lambda}_{\text{ML}}}, \tag{12}$$

which corresponds to the canonical definition of the temperature [Eq. (8)]. Because $\overline{f}$ is a sufficient statistic, it is not necessary to observe individual configurations to apply the formalism but also possible to work with ensemble data.

For infinitely many data ($N \to \infty$), the likelihood function becomes a δ peak and ML estimation is identical to the maximum entropy method [18], which also arrives at Eq. (12). The maximum entropy method derives the Gibbs ensemble [Eq. (2)] by searching for the distribution $p(x)$ that minimizes the relative entropy with respect to the reference distribution

$$H(p) = \int p(x) \ln[p(x)/\rho(x)]\, dx$$

and satisfies the constraints (12). The interactions λ become Lagrange parameters that have to be chosen such that the constraints (12) are satisfied. The maximum attainable entropy is [18]

$$H_{\max} = -\ln Z(\lambda) - \lambda\overline{f} = \tfrac{1}{N}\ln L(\lambda) \tag{13}$$

and proportional to the logarithm of the likelihood function. Therefore, the maximum likelihood and maximum entropy approaches are formally equivalent for the linear Hamiltonians [Eq. (1)] considered in this article. Conceptually, however, both approaches are different because maximum entropy *infers* the functional form of the Gibbs ensemble [Eq. (2)] and then estimates its parameters, whereas maximum likelihood *assumes* that the Hamiltonian [Eq. (1)] is correct.

B. Bayesian approach

A Bayesian approach extends maximum likelihood estimation by assigning a prior probability $p(\lambda)$ over the interaction parameters and considering their posterior probability:

$$p(\lambda|f_1,\ldots,f_N) \propto \left(\frac{e^{-\lambda\cdot\overline{f}}}{Z(\lambda)}\right)^N p(\lambda). \tag{14}$$

Maximization of $p(\lambda|\{f\})$ results in the maximum *a posteriori* (MAP) estimate. A more adequate representation of λ, especially its uncertainty, is achieved by drawing samples from the posterior distribution $p(\lambda|\{f\})$. Because

$$-\frac{\partial^2}{\partial\lambda_k\partial\lambda_l}\ln L(\lambda) = N\langle\delta f_k\delta f_l\rangle_\lambda$$

with $\delta f_k = f_k - \langle f_k\rangle_\lambda$ is positive semidefinite, the likelihood function is log concave but not necessarily in a strict sense. However, if we assign a strictly log-concave prior $p(\lambda)$, the posterior will exhibit a unique maximum. By considering an entire distribution of interaction parameters [Eq. (14)] we are estimating not only a single ensemble from the given data but a hyperensemble [21], i.e., an "ensemble of ensembles" that peaks at the maximum entropy ensemble. As the number of data or the system size grows, the ensembles will collapse into a single ensemble.

C. Sequential Monte Carlo algorithm

Although the posterior distribution of the interaction parameters [Eq. (14)] is unimodal, sampling from it is not straightforward due to the intractability of the partition function $Z(\lambda)$. Numerical integration or brute force enumeration of all configurations is only possible for small and finite systems. We outline a sequential Monte Carlo algorithm that estimates λ and $g(f)$ successively by constructing a sequence of approximations $p^{(t)}(\lambda|\{f\})$ and $g^{(t)}(f)$ where t is the iteration index. For given $\lambda^{(t)}$, M configurations $x_m^{(t)}$ are generated according to

$$x_m^{(t)} \sim p(x|\lambda = \lambda^{(t)}), \quad m = 1,\ldots,M, \tag{15}$$

where "$\sim$" means "sampled from." Using all the configurations sampled up to iteration t we estimate $g(f)$ by means of reweighting techniques [22,23]. In iteration t, the estimated feature distribution is

$$g^{(t)}(f) = \sum_{s=1}^{t}\sum_{m=1}^{M} g_{sm}^{(t)} \prod_k \delta\big(f_k - f_k\big(x_m^{(s)}\big)\big), \tag{16}$$

where $x_m^{(s)}$ is the mth configuration generated in iterations $s \leqslant t$ with feature vectors $f_m^{(s)} = f(x_m^{(s)})$ and weights $g_{sm}^{(t)}$, which are updated in every iteration. We generate new parameters $\lambda^{(t+1)}$ by drawing from

$$p^{(t)}(\lambda|\{f\}) \propto p(\lambda)\exp\left\{-N\left[\lambda\overline{f} + \ln\sum_{s,m} g_{sm}^{(t)} e^{-\lambda\cdot f_m^{(s)}}\right]\right\} \tag{17}$$

and start with $\lambda^{(0)} = 0$. In contrast to existing Monte Carlo approaches [9,24,25] the use of reweighting techniques allows us to incorporate all configurations generated during the Monte Carlo iterations and thereby improve the accuracy of the ensemble averages that are required to estimate λ [Eqs. (3) and (8)].

The approximate posterior $p^{(t)}(\lambda|\{f\})$ is also log-concave. One could determine its maximum by convex optimization or iterative scaling [26], a widespread method for learning maximum entropy models in natural language processing. Here we are interested in sampling the parameters to also explore their uncertainty. We use Hamiltonian Monte Carlo (HMC) [27] to draw λ from $p^{(t)}$ by starting from the previous sample $\lambda^{(t)}$ and using momenta generated from a standard normal distribution. During HMC we need to evaluate the forces stemming from the likelihood function $L(\lambda)$ [Eq. (11)]. These are proportional to $N(\langle f\rangle_\lambda - \overline{f})$ and drive the parameters so as to match the ensemble average and the observed average data.

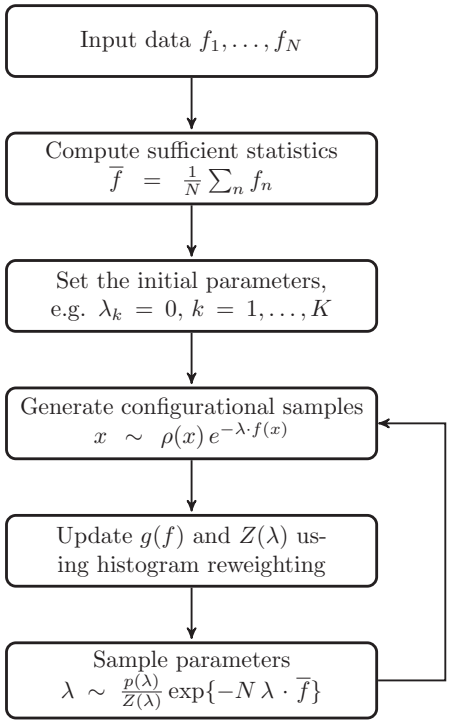

FIG. 1. Flow chart of the sequential Monte Carlo algorithm. Superscripts and indices have been omitted for clarity. Details of the individual steps are given in Eqs. (15), (16), and (17).

The flow chart (Fig. 1) gives an overview over the sequential Monte Carlo algorithm.

III. APPLICATIONS

A. Temperature estimation

To illustrate the algorithm we first estimate the temperature of a 16×16 Ising model from $N = 10$ energies generated at $\beta = 0.5$ with $\overline{E} = -446$. As outlined before we have a single descriptor $E(x) = -\sum_{\langle ij\rangle} x_i x_j$, and the unknown parameter is the inverse temperature β (the coupling is assumed known and fixed $J = 1$) to which we assign the exponential prior probability $p(\beta) = e^{-\beta}$. In every iteration ten configurations are generated ($M = 10$).

Figure 2(a) shows the final estimate of the feature distribution $g(E)$. The microcanonical estimator (10) demands that we accurately compute the feature distribution in the energy range that contains the observed average $\overline{E}$. Indeed the estimated feature distribution corresponds closely to the exact density of states [28] in the relevant energy range. Figure 2(b) shows the convergence of the approximate posterior distributions $p^{(t)}$ constructed during sequential Monte Carlo sampling. The color encodes the stage of the algorithm. In the initial phase

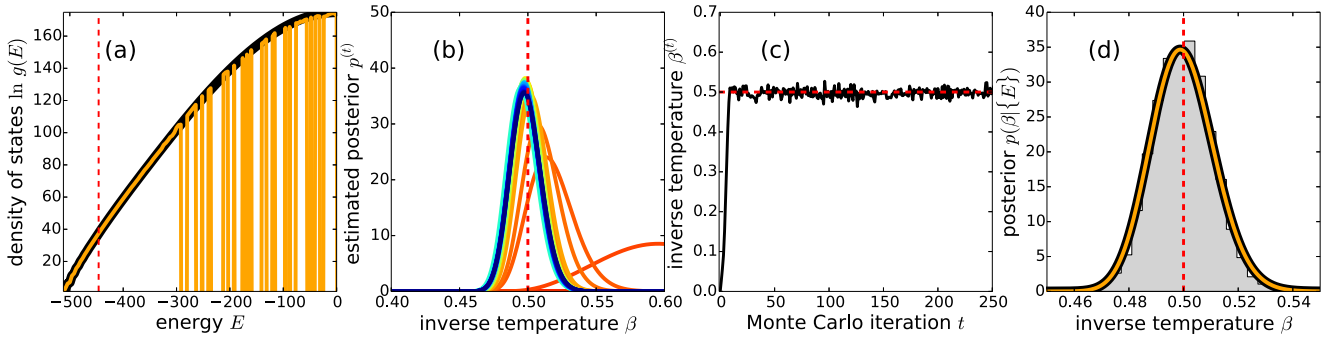

FIG. 2. (Color online) Estimating the temperature of the Ising model: (a) The estimated feature distribution is shown as orange (thin gray) line. The black (thick) curve is the true density of states. Shown are the natural logarithms of the feature distributions. The red (gray) dashed line indicates the mean of the 10 energies from which the temperature was estimated. (b) Estimated posterior distributions $p^{(t)}$. (c) Sampled inverse temperatures $\beta^{(t)}$. The dashed red (gray) line is the true inverse temperature. (d) Final posterior distribution of the temperature after convergence. Black (thick) line: true posterior distribution; orange (thin gray) curve: estimated posterior distribution based on the final estimate of the DOS. The gray histogram shows the sampled $\beta^{(t)}$.

(red to yellow to green distributions online) the posterior is biased towards too large β. Because the stepsize of the HMC sampler is set to a small value initially and then adjusted automatically, the algorithm has enough time to correct for the bias in $p(\lambda|\{E\})$. The approximate distributions converge to the correct posterior probability (blue [gray] distribution) in less than 50 iterations. Figures 2(c) and 2(d) show the first 250 temperature estimates and the final β posterior. Due to the small number of observed energies the posterior exhibits some degree of uncertainty with a mean and standard deviation of $\beta = 0.500 \pm 0.011$. The canonical and microcanonical temperature based on the same data are $\hat{\beta}_{\mathrm{ML}} = 0.499$ and $\hat{\beta}_{\mathrm{micro}} = 0.493$ (obtained by linear interpolation of the discrete density of states).

B. Inverse Ising problem

Let us now look at the inverse Ising problem, in which we aim to reconstruct the interaction parameters of the Sherrington-Kirkpatrick (SK) spin glass [29]:

$$E(x) = -\sum_{i<j} J_{ij} x_i x_j = \sum_k \lambda_k f_k(x), \tag{18}$$

where x is a D dimensional spin configuration, $x_i = \pm 1$, and J_{ij} are drawn from a Gaussian distribution with mean zero and variance D^{-1}. The dimension of the feature and parameter vectors, $f_k(x) = -x_i x_j$ and $\lambda_k = J_{ij}$, is $D(D-1)/2$ and exceeds the size of the configuration vector (the two configuration indices i, j are mapped to a single feature index k). But there are only 2^{D-1} unique feature vectors that are distributed uniformly. We compare the performance of the sequential Monte Carlo sampler with naive mean-field (nMF) theory [14] and pseudolikelihood maximization (PLM) [15], which was shown to provide highly accurate solutions of the inverse Ising problem. The quality of the reconstructed couplings is assessed in terms of the mean-squared error (MSE).

We reconstructed the interactions of a 20-dimensional SK model from 1000 configurations. For this system size the partition function and feature distribution can be evaluated by exhaustive enumeration. This allows us to compute the ML estimate based on the correct likelihood function [Eq. (11)] and compare it with the approximation found by the sequential Monte Carlo sampler. Figure 3(a) shows the evolution of the reconstruction error, which achieves a similar accuracy as exact ML based on the correct partition function. In the initial phase, the approximate log likelihood differs largely from the correct one. But as the sequential sampler proceeds, the approximation improves and converges to the correct log likelihood. The mean of the sampled interactions is more accurate than any of the individual $\lambda^{(t)}$ and performs similarly to nMF and PLM. Figure 3(b) shows the performance of the algorithms at different temperatures. The nMF approach has difficulties to obtain accurate reconstructions at low temperatures. PLM and sequential Monte Carlo both work for temperatures down to 0.5. But at lower temperatures PLM runs into problems whereas the Bayesian approach still works. An interesting property of the sequential Monte Carlo sampler is that at lower temperatures fewer configurations are needed to represent the feature distribution. The number of configurations at which $g(f)$ is estimated drops with increasing β and therefore the algorithm converges faster at lower temperatures.

C. Maximum entropy inference

As outlined in Sec. II A the maximum entropy method is formally equivalent to the Bayesian approach if the number of observations N approaches infinity. A simple modification of our algorithm can deal with this situation. Instead of sampling λ from $p^{(t)}$ [Eq. (17)], we set

$$\lambda^{(t+1)} = \arg\max_{\lambda} \ln p^{(t)}(\lambda|\{f\}). \tag{19}$$

Because this is a convex optimization problem it can be solved efficiently using optimizers such as the Powell algorithm.

To illustrate the application of our algorithm to maximum entropy inference, we apply our algorithm to the toy system studied by Roux and Weare [30]. This system has a single conformational degree of freedom x that follows the reference distribution

$$\rho(x) \propto \exp\left\{ -25\,(x-0.25)^4 + x\cos x - \frac{2\sin(20x)}{1+2x^2} \right\}.$$

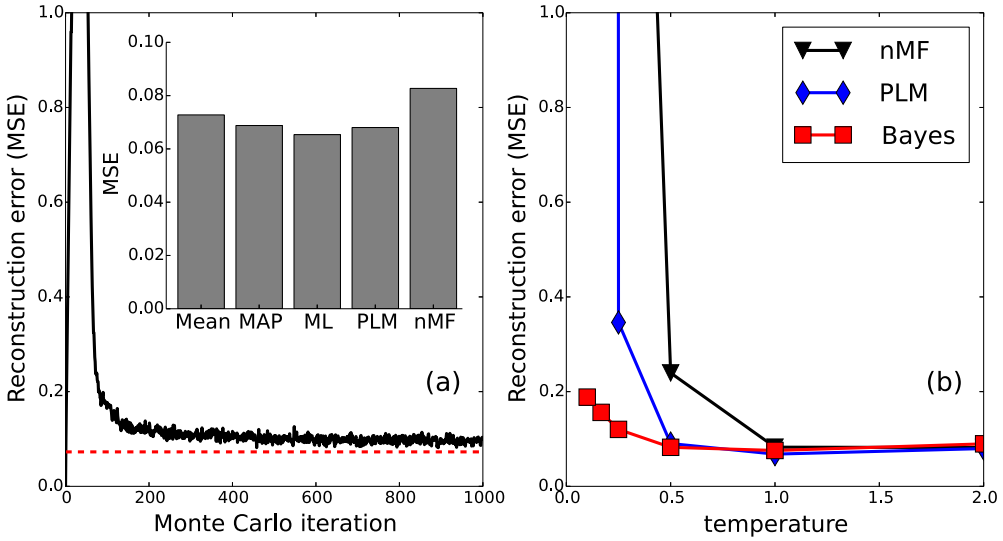

FIG. 3. (Color online) Inverse Ising problem. (a) Evolution of the reconstruction error (root mean-squared error, MSE) (black line). The dashed line indicates the MSE of the mean interaction parameters. The inset shows the performance of the Bayesian estimators (posterior mean and maximum) in comparison to exact ML, PLM, and nMF. (b) Inverse Ising problem at different temperatures.

Observed is the average position $\overline{x} = -0.127$, i.e., the only feature is $f(x) = x$. We are looking for the Lagrange parameter λ that reproduces the observed average position; the correct value is $\lambda = 10$. For given $\lambda^{(t)}$ new configurations are generated from $p(x|\lambda^{(t)}) \propto \rho(x)e^{-\lambda^{(t)}x}$ using the metropolis algorithm. Note that in contrast to recent applications of the maximum entropy method to biomolecular systems [30,31] our approach only requires simulations involving a *single* configuration and not an entire ensemble of configurations, which are coupled through an ensemble restraint.

Figure 4(a) shows the evolution of the mean position during the Monte Carlo iterations. The initial value is $\langle x \rangle = 0.258$ ($\lambda = 0$) and decreases towards the target value within 100 iterations. As the Monte Carlo procedure progresses the effective potential $-\ln \rho(x) + \lambda^{(t)} x$ approaches the target energy [Fig. 4(b)]. Figure 4(c) shows a histogram of the sampled configurations $x_m^{(t)}$. Although its main purpose is to generate samples of the maximum entropy parameters, the Monte Carlo sampler generates configurational samples that follow the correct distribution as a by-product.

D. Estimation of interaction potentials

Finally we apply the algorithm to recover the LJ parameters ϵ and σ from simulated configurations of a monatomic fluid. Several approaches including direct Boltzmann inversion [32], reverse and inverse Monte Carlo [9,24,25], potential and force matching [33,34], generalized Yvon-Born-Green approximation [8], as well as the configurational temperature formalism [10] have been used to estimate molecular inter-

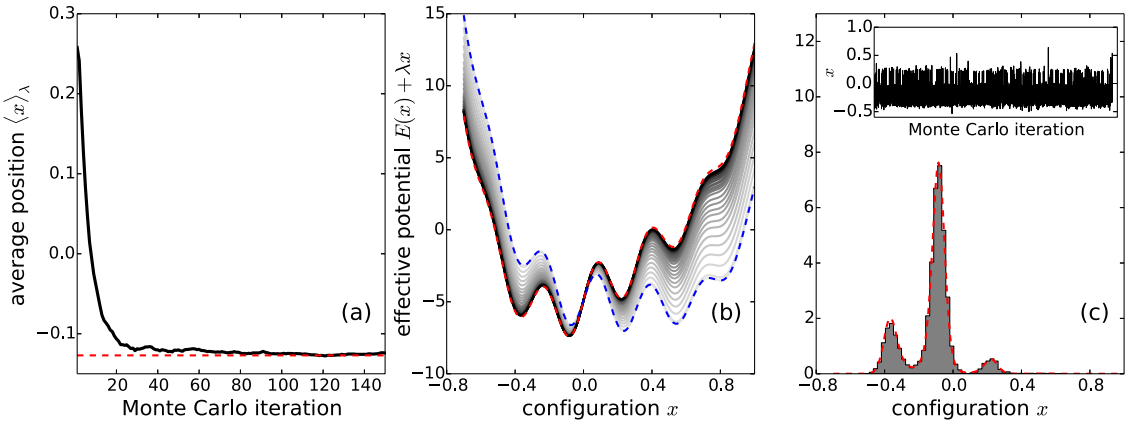

FIG. 4. (Color online) Application of the Monte Carlo algorithm to a maximum entropy fitting problem. The one-dimensional toy system [30] is analyzed using our sequential Monte Carlo sampler. (a) Evolution of the mean position $\langle x \rangle_{\lambda^{(t)}}$ (black line); the dashed red (gray) line indicates the observed average value -0.127. (b) Effective potential energy functions during Monte Carlo sampling. The light gray potential $E_0(x) = -\ln \rho(x)$ is gradually deformed by the addition of $\lambda\, x$. With increasing number of iterations the effective potential energy curves are shown by increasingly darker lines. Furthermore, the reference potential $E_0(x)$ and the target potential $E(x) = E_0(x) + 10\, x$ are highlighted by the dashed blue and red (gray) lines, respectively. (c) Histogram of the generated configurational samples; the inset shows the individual configurations generated after convergence. The correct maximum entropy distribution is shown as a red (gray) dashed line.

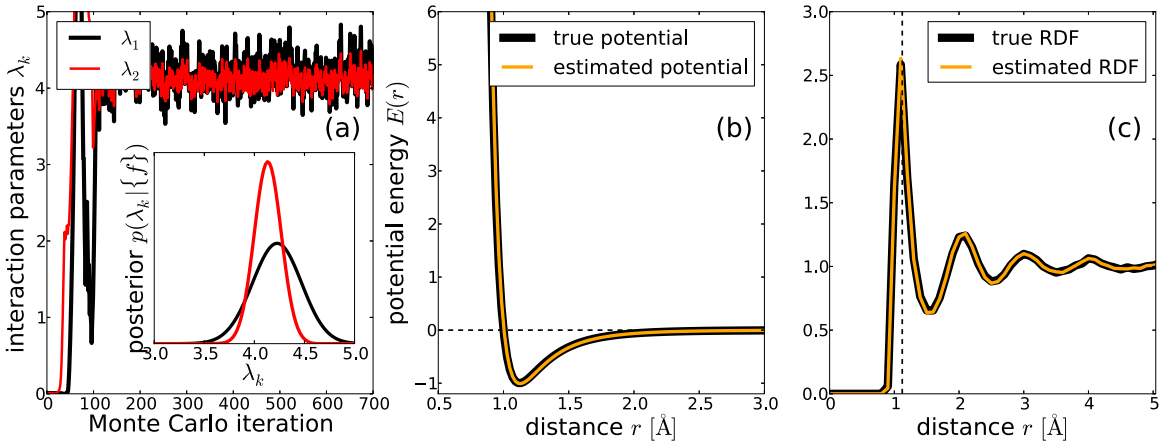

FIG. 5. (Color online) Estimation of a Lennard-Jones potential. (a) Sampled interaction parameters in the initial phase of the Monte Carlo algorithm. The inset shows the final estimated posterior distribution of the interaction parameters (λ_1 black, λ_2 light red [gray]). (b) Reconstructed and correct energy function shown as orange (thin gray) and black (thick) curve, respectively. (c) Radial distribution function (RDF) obtained with the proposed Monte Carlo sampler (thin orange [gray] line) compared to the correct RDF obtained with a replica-exchange simulation of the fluid using the true LJ parameters (thick black line).

action potentials. As outlined in the introduction there are two features corresponding to the repulsive and attractive part of the LJ potential. A system comprising 108 particles was simulated at $\beta = 1$ using HMC [27] and confined to a box of edge length 5.04 Å during the simulation. The parameters were set to $\epsilon = \sigma = 1$; therefore the true interaction parameters are $\lambda = (4,4)$. The interaction potential is recovered from the ensemble average $\overline{f} = (-320,170)$ calculated over 20 simulated configurations. Again the exponential distribution $p(\lambda_1,\lambda_2) = e^{-\lambda_1-\lambda_2}$ was used as prior probability. Within 100 iterations the sampler produces stable interaction parameters that scatter about the correct values: $\lambda_1 = 4.22 \pm 0.24$, $\lambda_2 = 4.13 \pm 0.13$ [Fig. 5(a)]. The microcanonical estimator (10) based on the estimated feature distribution gives $\hat{\lambda}_{\text{micro}} = (3.8,3.9)$. The estimated interaction potential accurately reproduces the correct potential used in the simulation [Fig. 5(b)]. The radial distribution function is often used to assess the quality of an estimated interaction potential. After convergence of the algorithm, the sampled configurations exhibit the correct radial distribution function [Fig. 5(c)].

IV. CONCLUSION

In summary, we formulate a Bayesian approach to solve inverse problems in statistical mechanics and propose a sequential Monte Carlo sampler that enables the application of the formalism to complex systems. On the inverse Ising problem the algorithm achieves accuracy similar to that of PLM but it is more flexible. First, we do not need to have access to the individual configurations but can work with average data directly. Second, our approach is also applicable to continuous systems for which it is not possible to compute the pseudolikelihood. In addition to existing methods such as nMF [14] and PLM [15] the Bayesian approach does not give a single point estimate but explores the full posterior distribution, thereby taking into account the uncertainty in the interaction parameters. For continuous systems the conformational temperature formalism [10] also requires that we have access to individual configurations of the system. Since the liklihood function only depends on the data through ensemble averages, the Monte Carlo algorithm is more broadly applicable and could also be used to analyze average data computed from different systems, as it is relevant for the derivation of "knowledge-based potentials" in biomolecular modeling.

ACKNOWLEDGMENT

This work has been supported by Deutsche Forschungsgemeinschaft (DFG) Grant No. HA 5918/1-1.

[1] W. Maysenhölder, Phys. Status Solidi B **139**, 399 (1987).
[2] E. Schneidman, M. J. Berry, R. Segev, and W. Bialek, Nature (London) **440**, 1007 (2006).
[3] T. R. Lezon, J. R. Banavar, M. Cieplak, A. Maritan, and N. V. Fedoroff, Proc. Natl. Acad. Sci. USA **103**, 19033 (2006).
[4] S. Cocco, S. Leibler, and R. Monasson, Proc. Natl. Acad. Sci. USA **106**, 14058 (2009).
[5] D. S. Marks, L. J. Colwell, R. Sheridan, T. A. Hopf, A. Pagnani, R. Zecchina, and C. Sander, PLoS ONE **6**, e28766 (2011).
[6] S. Torquato, Soft Matter **5**, 1157 (2009).
[7] H. Cohn and A. Kumar, Proc. Natl. Acad. Sci. USA **106**, 9570 (2009).
[8] J. W. Mullinax and W. G. Noid, Proc. Natl. Acad. Sci. USA **107**, 19867 (2010).
[9] A. Savelyev and G. A. Papoian, Proc. Natl. Acad. Sci. USA **107**, 20340 (2010).
[10] M. Mechelke and M. Habeck, J. Chem. Theor. Comput. **9**, 5685 (2013).
[11] B. Mandelbrot, Ann. Math. Stat. **33**, 1021 (1962).
[12] B. B. Mandelbrot, Phys. Today **42**, 71 (1989).

[13] M. Falcioni, D. Villamaina, A. Vulpiani, A. Puglisi, and A. Sarracino, Am. J. Phys. **79**, 777 (2011).
[14] H. J. Kappen and F. B. Rodriguéz, Neural Comput. **10**, 1137 (1998).
[15] E. Aurell and M. Ekeberg, Phys. Rev. Lett. **108**, 090201 (2012).
[16] M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, and T. Hwa, Proc. Natl. Acad. Sci. USA **106**, 67 (2009).
[17] J. Sohl-Dickstein, P. B. Battaglino, and M. R. DeWeese, Phys. Rev. Lett. **107**, 220601 (2011).
[18] E. T. Jaynes, Phys. Rev. **106**, 620 (1957).
[19] H. H. Rugh, Phys. Rev. Lett. **78**, 772 (1997).
[20] O. G. Jepps, G. Ayton, and D. J. Evans, Phys. Rev. E **62**, 4757 (2000).
[21] G. E. Crooks, Phys. Rev. E **75**, 041119 (2007).
[22] A. M. Ferrenberg and R. H. Swendsen, Phys. Rev. Lett. **63**, 1195 (1989).
[23] M. Habeck, Phys. Rev. Lett. **109**, 100601 (2012).
[24] R. McGreevy and L. Pusztai, Mol. Simul. **1**, 359 (1988).
[25] A. P. Lyubartsev and A. Laaksonen, Phys. Rev. E **52**, 3730 (1995).
[26] J. N. Darroch and D. Ratcliff, Ann. Math. Stat. **43**, 1470 (1972).
[27] S. Duane, A. D. Kennedy, B. Pendleton, and D. Roweth, Phys. Lett. B **195**, 216 (1987).
[28] P. D. Beale, Phys. Rev. Lett. **76**, 78 (1996).
[29] D. Sherrington and S. Kirkpatrick, Phys. Rev. Lett. **35**, 1792 (1975).
[30] B. Roux and J. Weare, J. Chem. Phys. **138**, 084107 (2013).
[31] A. Cavalli, C. Camilloni, and M. Vendruscolo, J. Chem. Phys. **138**, 094112 (2013).
[32] W. Tschöp, K. Kremer, J. Batoulis, T. Bürger, and O. Hahn, Acta Polym. **49**, 61 (1998).
[33] G. Tóth, J. Phys.: Condens. Matter **19**, 335222 (2007).
[34] F. Ercolessi and J. B. Adams, Europhys. Lett. **26**, 583 (1994).