Inferring hidden states in a random kinetic Ising model

Ludovica Bachschmid-Romano and Manfred Opper
Artificial Intelligence group, Computer Science, TU Berlin, Germany

Abstract

- We consider the problem of predicting the spin states in a kinetic Ising model when spin trajectories are observed for only a finite fraction of sites.
- Exact inference of hidden states is not tractable for large networks, but algorithms which are based on statistical physics approximations have recently been discussed [2, 3]. Hence, it will be interesting and important to study a scenario for which the theoretically optimal performance for predicting hidden spins can be computed exactly.
- We will show that such a solution can be found in the thermodynamic limit of an infinitely large network when the couplings are random. Our approach will be based on the replica method of disordered systems.
- We will then translate the picture derived for the disorder averaged system into equations for the local magnetization of hidden spins which are valid for a typical *single* system with fixed couplings and observations. As following step, such equations could be used to write an optimal algorithm for inferring the network couplings.

The model and Bayes optimal inference

We will consider a model with N Ising spins, divided into two groups:

- a group of observed spins $s_i(t)$, $i = 1 \dots N_{\rm obs} = \lambda N$, $t = 0 \dots T$,
- a group of hidden spins $\sigma_a(t)$ $a = 1 \dots N_{\rm hid} = (1-\lambda)N$, $t = 0 \dots T$

We assume parallel Markovian dynamics for the entire spin system, which is governed by the transition probability

$$P[\{s,\sigma\}(t+1)|\{s,\sigma\}(t)] = \prod_i \frac{e^{s_i(t+1)g_i(t)}}{2\cosh[g_i(t)]} \prod_a \frac{e^{\sigma_a(t+1)g_a(t)}}{2\cosh[g_a(t)]} \,,$$

where the fields are defined as

$$g_i(t) = \sum_j J_{ij}s_j(t) + \sum_b J_{ib}\sigma_b(t)\,, \qquad g_a(t) = \sum_j J_{aj}s_j(t) + \sum_b J_{ab}\sigma_b(t)\,,$$

in terms of the couplings J. In a Bayesian setting, where the probabilistic model of the spin dynamics is assumed to be known, the optimal prediction can be computed from the conditional (posterior) distribution of unobserved spins given the observed ones. Given a true 'teacher' sequence $\{\sigma^*\}$ of unobserved spins, we are interested in the total quality of the Bayes optimal prediction, given by the Bayes error

$$\varepsilon = \sum_{\{s,\sigma^*\}} P(\{s,\sigma^*\})\Theta(-\sigma^*_a(t)m_a(t)) = \sum_{\{s\}} P(\{s\}) \sum_{\{\sigma^*\}} P(\{\sigma^*\}|\{s\})\Theta(-\sigma^*_a(t)m_a(t)).$$

We will use the replica method to compute the error in the thermodynamic limit $N \to \infty$, when the couplings J are assumed to be mutually independent Gaussian random variables, with zero mean and variance k^2/N.

Replica analysis

The replica analysis reveales a fairly simple statistical picture of the posterior trajectories of hidden spins. Spins at different time steps (and sites) are *statistically independent*, but their local magnetizations depend on the propagation of information from past and future spins which is expressed through the order parameter

$$Q^{\alpha\beta}(t) = \frac{1}{N_{\rm hid}} \sum_a \sigma^\alpha_a(t)\sigma^\beta_a(t),\ \alpha \neq \beta$$

and its conjugate parameter. Assuming replica symmetry ($Q^{\alpha\beta}(t) = Q(t)\ \forall \alpha \neq \beta, t$), the selfaveraging values of the **order parameters** are:

$$Q(t) = \frac{1}{\langle W(t-1)\rangle_{\zeta_{t-1},\phi_t,\psi_{t-1}}} \left\langle \frac{\langle \tanh A(t-1) W(t-1)\rangle^2_{\zeta_{t-1}}}{\langle W(t-1)\rangle_{\zeta_{t-1}}} \right\rangle_{\phi_t,\psi_{t-1}} , \qquad t = 1 \dots T$$

$$\hat{Q}(t) = \frac{ik^2(1-\lambda)}{\langle W(t)\rangle_{\zeta_t,\phi_{t+1},\psi_t}} \left\langle \frac{\langle [\tanh A(t) - \tanh B(t)] W(t)\rangle^2_{\zeta_t}}{\langle W(t)\rangle_{\zeta_t}} \right\rangle_{\phi_{t+1},\psi_t}$$

$$+\, ik^2\lambda \sum_{\{s\}(t+1)} \left\langle \frac{\langle [s(t+1) - \tanh B(t)] V(t)\rangle^2_{\zeta_t}}{\langle V(t)\rangle_{\zeta_t}} \right\rangle_{\psi_t} , \qquad t = 0 \dots T-1$$

where

$$A(t) = \psi(t) + \zeta(t) + \phi(t+1), \qquad B(t) = \psi(t) + \zeta(t),$$

$$W(t) = \frac{\cosh A(t)}{\cosh B(t)}, \qquad V(t) = \frac{e^{s(t+1)B(t)}}{2\cosh B(t)}$$

and $\zeta(t), \psi(t), \phi(t)$ are Gaussian independent random fields with zero mean and covariances:

$$\langle \psi(t)\psi(t)\rangle = k^2\left(\lambda + (1-\lambda)Q(t)\right), \qquad \langle \zeta(t)\zeta(t)\rangle = k^2(1-\lambda)\left(1 - Q(t)\right),$$

$$\langle \phi(t)\phi(t)\rangle = -i\hat{Q}(t).$$

Distribution of local magnetization: $p_t(m) = \langle w(\psi,\phi)\delta(m - m(t|\psi,\phi)\rangle_{\psi_{t-1},\phi_t}$, where

$$m(t|\psi,\phi) = \frac{\langle \tanh A(t-1) W(t-1)\rangle_{\zeta_{t-1}}}{\langle W(t-1)\rangle_{\zeta_{t-1}}}, \qquad w(\psi,\phi) = \frac{\langle W(t-1)\rangle_{\zeta_{t-1}}}{\langle W(t-1)\rangle_{\zeta_{t-1},\psi_{t-1},\phi_t}}.$$

The **Bayes error** equation is translated into: $\varepsilon = \frac{1}{2}\left(1 - \int_{-1}^{1} p_t(m)\,|m|\,dm\right)$.

Results

On the *right*: order parameter Q as a function of time.

Our analytical results agree very well with simulations of spin systems with relative small number of spins. For these systems we could compute local magnetizations $m_a(t)$ exactly by enumeration. The Markovian spin dynamics facilitated these computations with the use of a *forward–backward* algorithm. We then compute Q using $Q(t) = \frac{1}{N_{\rm hid}} \sum_{a=1}^{N_{\rm hid}} E_{s,J}\, m_a^2(t)$.

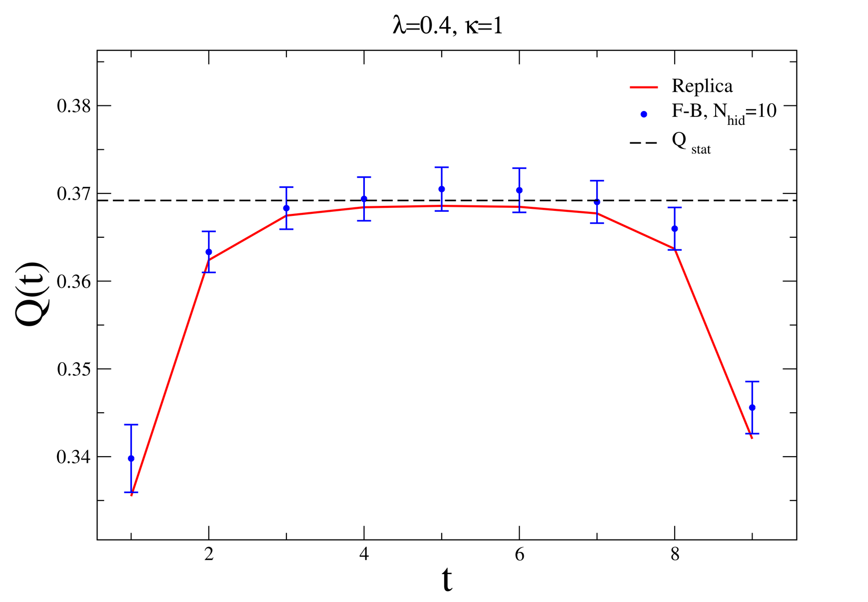

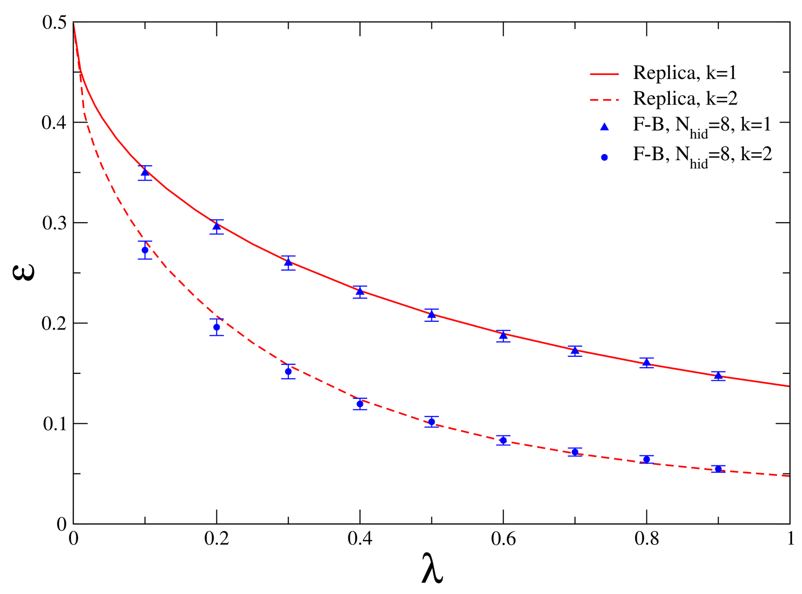

On the *left*: Bayes error as a function of the load factor λ.

In the limit of no observations, the prediction on the the state of hidden spins is completely random ($\varepsilon = 0.5$). The error rapidly decreases as λ gets larger, but remains nonzero for $\lambda = 1$, indicating the presence of a residual error in almost fully observed systems due to the stochasticity of the Markov process. Since the couplings are responsible for the propagation of information between spin sites, the Bayes error decreases as the coupling strength increases.

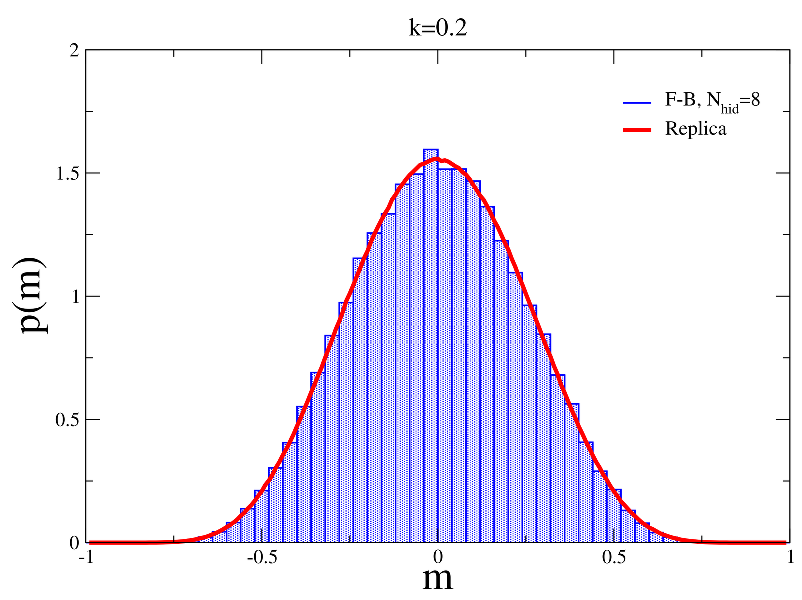

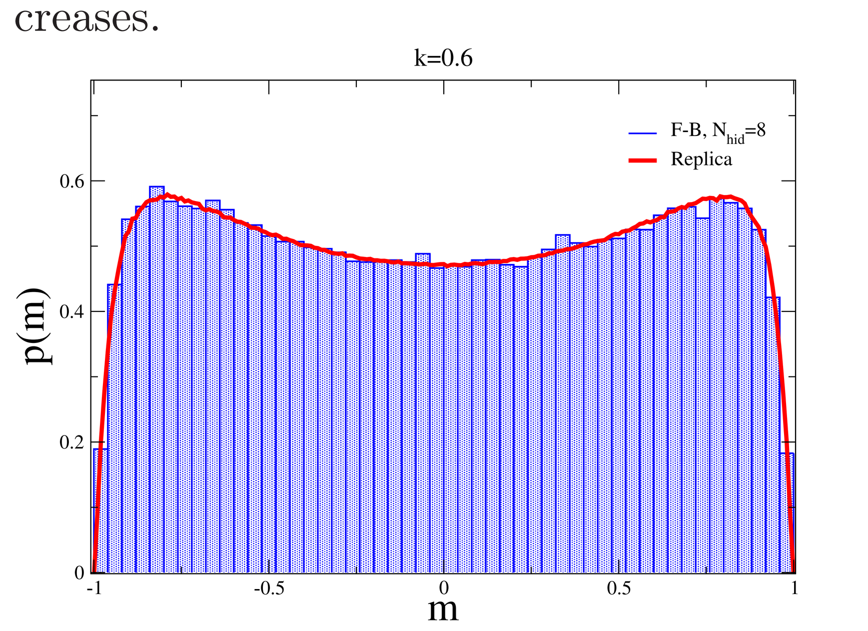

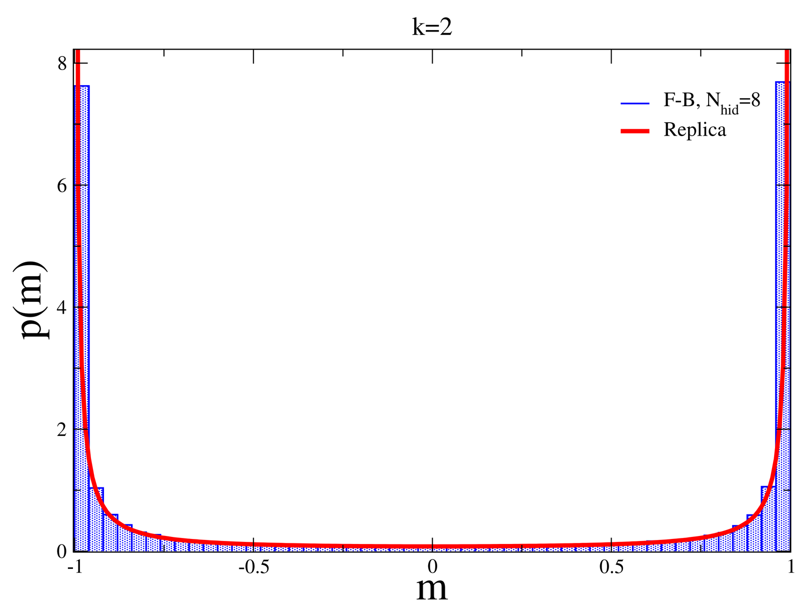

Distribution of local magnetization of hidden spins.

For small k the distribution is close to a Gaussian centered at zero, with vanishing variance as $k \to 0$, meaning that nontrivial prediction on the magnetization can be made. As k grows larger the distribution broadens and above a critical value the curve becomes bimodal. For large k, the distribution $p(m)$ concentrates at $m = \pm 1$, allowing for a perfect prediction of hidden spins.

Cavity approach

From cavity arguments we derive the equations for the local magnetization of hidden spins for a typical *single* system with fixed couplings and observations. Such mean field equations provide an efficient algorithm for the computation of local magnetizations in large random networks. This could then be used as an approximation in the E-Step of an EM algorithm [4] which aims at computing the maximum likelihood estimator of the network couplings, averaging out unobserved spins.

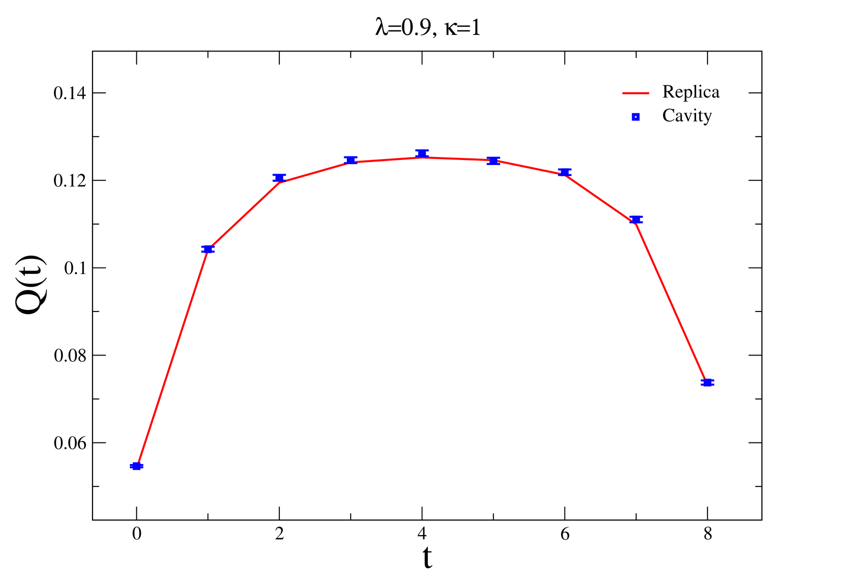

Order parameter Q as a function of time. In the cavity approach, we compute Q from the local magnetizations using $Q(t) = \frac{1}{N_{\rm hid}} \sum_{a=1}^{N_{\rm hid}} E_{s,J}\, m_a^2(t)$.

References

[1] L. Bachschmid-Romano and M. Opper. *J. Stat. Mech: Theor. Exp*, 2014(6):P06013,2014.
[2] B. Dunn and Y. Roudi.*Phys. Rev. E*, 87:022127, Feb 2013.
[3] Joanna Tyrcha and John Hertz.*MBE*, 11:149, Feb 2014.
[4] A. P. Dempster, N. M. Laird, and D. B. Rubin. *J. R. Stat. Soc. Series B*, 39:1-38, Dec 1977.