

The appropriateness of ignorance in the inverse kinetic Ising model

Benjamin Dunn, Claudia Battistin

Kavli Institute for Systems Neuroscience and Centre for Neural Computation, NTNU

Email: benjamin.dunn@ntnu.no, claudia.battistin@ntnu.no

Abstract

We develop efficient ways to consider and correct for the effects of hidden units for the paradigmatic case of the inverse kinetic Ising model with fully asymmetric couplings. We identify two sources of error in reconstructing the connectivity among the observed units while ignoring part of the network. One leads to a systematic bias in the inferred parameters, whereas the other involves correlations between the visible and hidden populations and has a magnitude that depends on the coupling strength. We estimate these two terms using a mean field approach and derive self-consistent equations for the couplings accounting for the systematic bias. Through application of these methods on simple networks of varying relative population size and connectivity strength, we assess how and under what conditions the hidden portion can influence inference and to what degree it can be crudely estimated. We find that for weak to moderately coupled systems, the effects of the hidden units is a simple rotation that can be easily corrected for. For strongly coupled systems, the non-systematic term becomes large and can no longer be safely ignored, further highlighting the importance of understanding the average strength of couplings for a given system of interest.

1 Introduction

Recent technological advances in high-throughput recordings of biological systems are enabling the use of statistical tools to ask an entirely new set of questions. The extent to which a given system can be observed as a whole, however, often limits the applicability of these tools. For example, currently the most impressive neural recordings are of populations of hundreds or thousands of neurons. While a vast improvement from the single recordings from approximately 60 years ago [30], it is a minuscule fraction of what is likely the relevant population. Still, with the hope of gaining insight into possible structure of the underlying network, scientists have turned to statistical models, such as the generalized linear model (GLM) [21], that allow potential *connectivity* (the structure of pairwise interactions and their strength) between the observed units to be estimated systematically. Of the many flavors of GLM, the simplest takes the form of a binary spin model or Bernoulli GLM [21], also known as the kinetic Ising model [12, 23] in statistical physics. This popular model takes the form,

$$p(s_i(t+1)|\mathbf{s}(t)) = \frac{\exp(s_i(t+1)H_i(t))}{2 \cosh(H_i(t))} \quad (1)$$

where $s_i(t) = \pm 1$ with $i = 1, \dots, N$ is the state of node i at time t . $H_i(t) = \sum_j J_{ij}s_j(t)$ is the input to node i , consisting of the field generated by the other nodes, with J_{ij} the coupling from node j to i . It should be noted that at this stage, following (1), we do not allow for an external field. Even if this model is typically regarded as a model of spiking neurons [6], since its introduction [17], it has been used for modeling financial markets [36, 4]. This model is indeed prone to applications in a variety of fields, from biology (e.g. gene interaction networks) to social sciences, where interactions are not symmetric and external inputs are non-stationary. Ultimately the apparent simplicity of the dynamical rule (1) along with its rich dynamics makes this framework ideally suited for developing techniques and gaining insight into the more general class of statistical modeling, and therefore it has recently attracted considerable attention from the statistical physics community [24, 20, 18, 34, 27, 25, 11].

When the system is fully observed, the learning of the parameters of this model can be accomplished using techniques that take advantage of the convexity of the problem [10]. These methods, however, are iterative in nature and often require calculating inverses of large matrices with entries consisting of averages over the data that have to be recomputed at each iteration. While not a problem for many data sets currently available, the promise of orders of magnitude more data would suggest that it will soon be necessary to turn to approximate methods. *Mean field* approximations [25, 24, 20, 18] that have been developed for this purpose, model the fluctuating field $H_i(t)$ with an effective field, reducing a many body problem to many

single body problems. We will focus on fully asymmetric connectivities for which mean field equations, exact in the thermodynamic limit, have been first introduced in [20]. Here by fully asymmetric couplings we mean fully asymmetric Sherrington-Kirkpatrick couplings [29]: the J_{ij} s are assumed to be independently drawn (J_{ij} statistically independent from J_{ji}) from a normal distribution with mean J_1/N and standard deviation J_0/\sqrt{N} . For large asymmetric networks the authors in [20] approach the inverse problem deriving a relation between correlations and couplings:

$$\mathbf{D}(t) = \mathbf{A}(t)\mathbf{J}\mathbf{C}(t) \quad (2)$$

where $D(t)/C(t)$ are the time-delayed and equal-time correlation matrices

$$D_{ij}(t) = \langle ds_i(t)ds_j(t-1) \rangle, \quad (3)$$

$$C_{ij}(t) = \langle ds_i(t)ds_j(t) \rangle. \quad (4)$$

Here $\langle \cdot \rangle$ indicates averaging over the distribution $P(\mathbf{s}) = \prod_{i,t} p(s_i(t+1)|\mathbf{s}(t))$, where $p(s_i(t+1)|\mathbf{s}(t))$ is defined in (1), while $ds_i(t) \equiv s_i(t) - m_i(t)$, where $m_i(t) = \langle s_i(t) \rangle$. In (2) the diagonal matrix \mathbf{A} is:

$$A_{ij}(t) = \delta_{ij} \int Dx [1 - \tanh^2 (g_i(t) + x\sqrt{\Delta_i(t)})] \quad (5)$$

where $\Delta_i(t) = \sum_k J_{ik}^2 (1 - m_k(t)^2)$, $Dx = \frac{dx}{\sqrt{2\pi}} e^{-x^2/2}$ and $g_i(t) = \langle H_i(t) \rangle$.

Interestingly, this relationship has been shown to be exact for fully observed systems with asymmetric couplings, even in the strongly coupled regime [20]. It is, however, very uncommon that a system is able to be fully observed, necessitating further understanding of what kind of errors we incur by ignoring the hidden population and perhaps develop approximations to account for them more explicitly. By explaining away correlations due to these unobserved but potentially relevant features, we begin to address the concerns we share with many about how informative are observed data and perhaps, some day, make stronger claims about our abilities to uncover actual network features.

Besides its potential applications, the hidden population problem in the KI model is intriguing both from computational and “physical” points of view. Computationally, since exact and Monte Carlo sampling of the states space are expensive if not unfeasible; physically, since the likelihood of the data in partially observed systems, as opposed to the fully observed ones, is not convex. For GLMs this problem has been approached previously by, for example, assuming the hidden units provide a source of Gaussian random noise, as in [15], and by approximate *Expectation Maximization* (EM) algorithms as in [22, 9, 33, 2]; see also [1]. In this framework [31, 7], the likelihood of the observed time series is maximized by alternatively estimating expected values of the joint (observed + hidden) distribution at fixed connectivity and then updating the couplings accordingly. When the hidden population constitutes a minor component of the entire network the EM algorithm has been proven accurate [35, 19], while in the highly subsampled regime the problem becomes burdened with local optima.

There are few cases where even in the presence of many hidden units, the connected and disconnected pairs, can be well identified from the inferred connections. For instance, [16] reported that increasing the number of genes used in their inference did not significantly change the important couplings between the observed nodes inferred using a maximum entropy model. In this specific case the insensitivity of the reconstruction to the inclusion of further nodes seems to follow from the specific choice of the inclusion criterion, such that the presence of the additional nodes does not affect significantly the correlations between the original ones. A second example regards a subpopulation of neurons: in [26] the connectivity exhibits a rescaling but no remarkable restructuring by observing larger and larger populations. This effect has been observed even when the generative model is an equilibrium model, as in [13] and can be attributed to the mean field properties of the underlying network, in a way that the unobserved neurons affect the correlations among the observed neurons as a whole. These results suggest that, when mean field arguments apply, the inference of the couplings between observed nodes can be corrected for the unobserved through a global transformation. In figure 1 A-B we show that the scaling of the couplings is not an effect limited to equilibrium models but it extends to the kinetic model with asymmetric couplings and does not rely on approximate learning.

In this paper we follow up on these ideas analyzing the effects of the hidden variables on the reconstruction of the couplings when you only see P neurons (\mathbf{J}^P), comparing them with those inferred from the full network

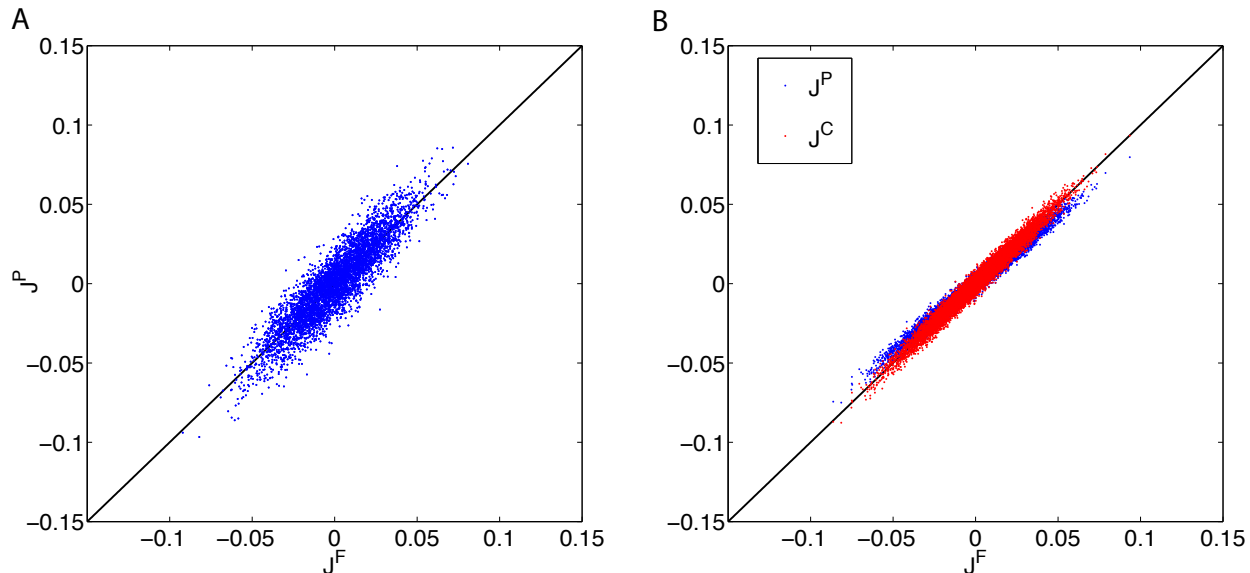


Figure 1: Network reconstruction in presence of hidden nodes. Scatter plot of the inferred couplings among $P = 100$ nodes out of $N = 500$: observing only the P nodes \mathbf{J}^P versus using data from the whole network \mathbf{J}^F . SK couplings \mathbf{J} independently drawn from a normal distribution with zero mean and standard deviation J_0/\sqrt{N} with $J_0 = 0.5$; data length $L = 10^7$ [†]. **A**: equilibrium Ising model, symmetrized version of the \mathbf{J} , inference performed using the Naïve Mean Field algorithm [32][14]. **B**: kinetic Ising model, fully asymmetric couplings \mathbf{J} ; inference performed using the mean field learning from [20] (blue)—not visually distinguishable from those learned via exact learning [12]—, corrected using (21) (red).

(\mathbf{J}^F), that we will partition as follows:

$$\mathbf{J}^F = \begin{bmatrix} \mathbf{J}_{ss}^F & \mathbf{J}_{s\sigma}^F \\ \mathbf{J}_{\sigma s}^F & \mathbf{J}_{\sigma\sigma}^F \end{bmatrix} \quad (6)$$

where s indicates the observed nodes, while σ the hidden ones. Exploiting (2) the difference between the inferred couplings can be written as:

$$\mathbf{J}^P - \mathbf{J}_{ss}^F = \mathbf{\Xi} + \mathbf{\Gamma} \quad (7)$$

with

$$\mathbf{\Xi} \equiv \left(\mathbf{I} - [\mathbf{A}_{ss}^F]^{-1} \mathbf{A}^P \right) \mathbf{J}^P \quad (8)$$

$$\mathbf{\Gamma} \equiv \mathbf{J}_{s\sigma}^F \mathbf{C}_{\sigma s} [\mathbf{C}_{ss}]^{-1} \quad (9)$$

where the matrix \mathbf{A} inherits the label, \mathbf{F} or \mathbf{P} , from the couplings that enter its expression and \mathbf{A}^F is partitioned, just like \mathbf{C} , accordingly to (6). For sake of clarity from (7) on we drop the time indices, bearing in mind that the inference problem can be solved in parallel at each time step. The following derivation will involve dynamical quantities at fixed time t , such that, for example, m_i will now indicate $m_i(t)$ instead.

Thus the error can be seen as having two components, a systematic term $\mathbf{\Xi}$ and a non-systematic one $\mathbf{\Gamma}$ that is shaped by the observed to observed and hidden to observed correlations.

In the next sections we seek to find a simple way to deal with the hidden population by estimating these two contributions in a mean-field type approach.

[†]In the regime of large L and in absence of an external field, we consider the process as being stationary, such that, when solving (2) we replace population averages $\langle \dots \rangle$ with time averages.

2 The systematic error term Ξ

In this section we are going to study the error term Ξ in (7). It generates a systematic bias in the learned couplings, that affects concertedly their magnitude, as shown in Figure 1.

We will now evaluate Ξ using a mean field approximation, building up on the assumption that statistically the hidden part of the network differs from the observed one neither architecturally nor dynamically. The local field on the observed node i can then be expressed as the mean field generated by the observed nodes plus a random variable ρ :

$$H_i = \sum_j^P J_{ij} m_j + \rho_i \quad (10)$$

For zero mean couplings ($J_1 = 0$) and fields, we have $\langle \rho_i \rangle = 0$ and we approximate its variance with

$$\text{var}(\rho_i) \approx \sum_k^P (J_{ik})^2 (1 - m_k^2) + (N - P) \text{var}(\mathbf{J}) \left(1 - \overline{\mathbf{m}^2}\right). \quad (11)$$

where $\bar{\cdot}$ and $\text{var}(\cdot)$ indicate respectively mean and variance over the space of the observed variables, namely $\overline{\mathbf{m}^2} \equiv \frac{1}{P} \sum_j m_j^2$ and $\text{var}(\mathbf{J}) = \frac{1}{P^2} \sum_{ij} (J_{ij} - \bar{\mathbf{J}})^2$. Notice that a non-zero mean of the couplings (e.g. J_1/N) would only affect the mean of the variable ρ , resulting into an additional term that, in the mean field framework can be approximated as $(N - P)\bar{\mathbf{J}}\bar{\mathbf{m}}$.

In the large systems limit, being the couplings J_{ij} iid random variables of order $1/\sqrt{N}$, the Central Limit theorem applies to the fields $H_i(t)$. Thus ρ is a Gaussian distributed random variable, such that A_{ij}^F in (7) becomes

$$\hat{A}_{ij}^F = \delta_{ij} \int Dx \left[1 - \tanh^2 \left(\sum_{k=1}^P J_{ik}^F m_k + x(\text{var}(\rho_i^F))^{\frac{1}{2}} \right) \right] \quad (12)$$

with $\text{var}(\rho_i^F)$ being the variance in (11) with couplings \mathbf{J}_{ss}^F inferred using the full data set.

Finally we can substitute (12) in the expression for the systematic component of the error:

$$\Xi \approx (\mathbf{I} - [\hat{\mathbf{A}}_{\text{ss}}^F]^{-1} \mathbf{A}^P) \mathbf{J}^P. \quad (13)$$

From (7) we can construct a set of self-consistent equation for the couplings that would be inferred using the full data set, $\hat{\mathbf{J}}_{\text{ss}}^F$, accounting for this systematic error on the couplings inferred from the reduced population, thus

$$\hat{\mathbf{J}}_{\text{ss}}^F = [\hat{\mathbf{A}}_{\text{ss}}^F]^{-1} \mathbf{A}^P \mathbf{J}^P - \Gamma \quad (14)$$

where in $\hat{\mathbf{A}}_{\text{ss}}^F$ we have replaced \mathbf{J}_{ss}^F with $\hat{\mathbf{J}}_{\text{ss}}^F$.

Within the limits of applicability of the mean field approximation developed in this section, we can now explain why the global structure of the weights is retained when ignoring the systematic error term in the reconstruction. Despite (14) is not an explicit relation between $\hat{\mathbf{J}}_{\text{ss}}^F$ and \mathbf{J}^P , we expect $\hat{\mathbf{A}}_{\text{ss}}^F$ to be weakly sensitive to the identity of the couplings that enter (12), when both couplings and states have zero mean. If this holds then the main difference between $\hat{\mathbf{A}}_{\text{ss}}^F$ and \mathbf{A}^P in (14), consist of a uniform shift in the variance of the gaussian integral associated with (12) given by the second term in (11). This produces a global rescaling of the couplings reflected in the rotation exemplified in Figure 1C.

We can now turn our attention to the second error term.

3 The non-systematic error term Γ

The component Γ of the error in (7) can be interpreted as a propagation of the correlations between observed and hidden variables on the correlations between observed units via the hidden to observed interactions. Notice that this term indeed vanishes when there is no hidden to observed coupling, even though the two populations might covary.

For couplings with a mean and standard deviation of zero ($J_1 = 0$) and J_0/\sqrt{N} , respectively, we have average $\bar{\mathbf{\Gamma}} = 0$. Thus in order to quantify the extent to which this error term affects the learned couplings one has to estimate its variance $\text{var}(\mathbf{\Gamma})$. Since we note that for small J_0 , the terms on the diagonal of $[\mathbf{C}_{\text{ss}}]^{-1}$ dominate

$$\text{var}(\mathbf{\Gamma}) \approx \frac{N-P}{N} J_0^2 \text{var}(\mathbf{C}^{\text{off}}) \quad (15)$$

where \mathbf{C}^{off} are the off-diagonal elements of \mathbf{C} . We can approximate \mathbf{C}^{off} by observing, as in [20], that the joint distribution of the fluctuations around the mean local field to each spin $\delta g_i = \sum_l J_{il} \delta s_l$ is $\epsilon_{ij} = \langle \delta g_i \delta g_j \rangle = [\mathbf{J} \mathbf{C} \mathbf{J}^T]_{ij}$. This leads to a joint distribution of $x = \delta g_i$ and $y = \delta g_j$ of the form [20]

$$P(x, y) = \frac{1}{2\pi \sqrt{\Delta_i \Delta_j}} \exp\left(-\frac{x^2}{2\Delta_i} - \frac{y^2}{2\Delta_j} + \epsilon_{ij} \frac{xy}{\Delta_i \Delta_j}\right) \quad (16)$$

Then equal-time correlations (3) can be written as:

$$\begin{aligned} C_{ij} &= \int dx dy P(x, y) \delta \tanh(h_i + g_i + x) \delta \tanh(h_j + g_j + y) \\ &\approx \frac{\epsilon_{ij}}{\Delta_i \Delta_j} \int \frac{dx}{\sqrt{2\pi \Delta_i}} \frac{dy}{\sqrt{2\pi \Delta_j}} \exp\left[-\frac{x^2}{2\Delta_i} - \frac{y^2}{2\Delta_j}\right] xy \delta \tanh(h_i + g_i + x) \delta \tanh(h_j + g_j + y) \\ &= \epsilon_{ij} a_i a_j \end{aligned} \quad (17)$$

where $A_{ij} = \delta_{ij} a_i$ is defined in (5). In (17) we have first expanded the exponential in the first line to the linear order in ϵ (recall that $\epsilon \propto N^{-1/2}$), then we performed the integration by parts to get to the third line. Using (17) one can show that the variance of the off-diagonal elements of \mathbf{C} is modulated by the couplings strength J_0

$$\text{var}(\mathbf{C}^{\text{off}}) \approx \frac{J_0^4 \text{var}(\mathbf{a})^2 \sum_j (1 - m_j^2)^2}{N^2 [1 - J_0^4 \text{var}(\mathbf{a})^2]} \quad (18)$$

where \mathbf{a} is the vector of the diagonal entries of \mathbf{A} defined in (5).

Replacing (15) in (18) we obtain an estimate for the variance of the non-systematic component of the error in (7):

$$\text{var}(\mathbf{\Gamma}) \approx \frac{J_0^6 (N-P) \text{var}(\mathbf{a})^2 \sum_j (1 - m_j^2)^2}{N^3 [1 - J_0^4 \text{var}(\mathbf{a})^2]} \quad (19)$$

$$\approx J_0^6 (N-P) / N^2 \quad (20)$$

where in the second line we note that, for small values of J_0 , one has $A_{ij} \approx \delta_{ij} (1 - m_i^2)$. From (20) one can identify two limiting behaviors of the non-systematic error term when compared with values of the inferred couplings, for large systems with weak interactions. If the fraction of hidden units is large ($\frac{N-P}{N} \sim O(1)$), the relative error $\text{var}(\mathbf{\Gamma})/\text{var}(\mathbf{J}) \sim J_0^4$ stays finite even in the regime of perfect sampling; while it vanishes if the fraction of hidden units is small. Notice that a non-zero mean of the couplings (e.g. J_1/N) won't change our estimate of $\text{var}(\mathbf{\Gamma})$ to the leading order.

Numerical simulations (Figure 2) confirm that for weak couplings, $J_0 = 0.1$ (Figure 2A) equation (20) provides a very accurate estimate of this error term while for stronger couplings $J_0 = 0.5$ (Figure 2B) this has already diverged. The approximation was not included for $J_0 = 1$ given that it lies outside the bounds of the figure.

Finally, for small values of J_0 , we can assume the non-systematic term $\mathbf{\Gamma}$ term is negligible and approximate the couplings using the self-consistent equation

$$\mathbf{J}^{\mathbf{C}} = [\mathbf{A}^{\mathbf{C}}]^{-1} \mathbf{A}^{\mathbf{P}} \mathbf{J}^{\mathbf{P}} \quad (21)$$

where $\mathbf{A}^{\mathbf{C}}$ is defined as $\hat{\mathbf{A}}^{\mathbf{F}}$ in (12) but replacing $\mathbf{J}^{\mathbf{F}}$ with $\mathbf{J}^{\mathbf{C}}$. Figure 1C shows an example of the correction to the inferred couplings from the full data set provided by (21), while Figure 2 demonstrates that using

(21) improves the learning systematically reducing the error in some cases to nearly half. Interestingly, the variance of the residual error, $\text{var}(\mathbf{J}_{\text{ss}}^{\mathbf{F}} - \mathbf{J}^{\mathbf{C}})$, matches well the variance of the non-systematic error term, $\mathbf{\Gamma}$ in (9) even for stronger values of J_0 and only diverging for the systems of both strong couplings ($J_0 = 1$) and large hidden populations (fraction visible < 0.5). This trend can be seen also in Figure 3 where the relative errors have been plotted, normalized by the variance of the couplings for systems of varying coupling strength and fraction visible. Here the regime where the correction can no longer improve the inference (top left corner) appears as a phase transition in the space of coupling strength and fraction observed. Importantly, in the regime where the correction can be useful, the error due to subsampling quickly becomes significant, as illustrated in Figure 4, surpassing the expected error due to data length limitations at reasonably small data lengths.

4 Discussion

In this paper we deal with the effects of couplings and states of the hidden variables in a mean-field type approximation on the inference of the couplings between the observed nodes. We showed that the hidden population provides two sources of errors. The first term can be largely attributed to the underestimated size of the system, while the second term takes into account the hidden to observed correlations and couplings. Neither of the two contributions to the error (7) depends explicitly on interactions between hidden units that introduce higher order corrections to the learning. The reconstruction of the visible part of the network is indeed basically unaffected when removing the hidden to hidden connections [9], motivating the development of algorithms that neglect them [2][33].

We showed that in case of weakly interacting units, the non-systematic error from ignoring the hidden population is small with respect to the strength of the couplings. In the weak couplings regime the non-systematic term can then be safely neglected and the systematic one treated using a mean field approximation. This allowed us to integrate the latter in the inference (equation (21)), resulting in a significantly better estimate of the true couplings, even for more strongly coupled systems. For networks with strong couplings, however, and, in particular, those with comparatively large hidden populations, the non-systematic error component dominates and could not be well addressed following this approach.

As mentioned earlier, in many data sets we are forced to ignore the large majority of the relevant network. Attempting to make conclusions from parameters inferred from such a tiny fraction seems ridiculous but the results here would suggest that, at least for weakly coupled systems, there might be hope. The question then begs itself, *how strongly connected are the systems we care about?* If strongly connected, as very well could be the case [28], then we may have reason to be concerned.

An alternative strategy to the reconstruction of partially observed networks, that also avoids addressing the dark side directly, would be to learn an external field in addition to the connectivity [3, 5, 36]. The external field would then *explain away* correlations between observed nodes driven by hidden nodes, preventing the observed to observed couplings to adjust for those. In this paper we intentionally do not include the external field in the model (1), even though the dynamical model allows for it. We indeed expect that, to some extent, the inferred external field will take care of the input from hidden to observed units, reducing the error in the reconstructed connectivity [15]. As an example, by learning time-varying fields that explain away correlations in the data due to aspects of the experiment, the authors in [8] were able to infer seemingly significant functional connectivity from a very small subpopulation of neurons. Further work should shed light on the extent to which the effects of the hidden population can be accounted for by including different non-stationary external fields.

Acknowledgments

We thank Yasser Roudi for his substantial contribution to the original ideas behind this manuscript and useful discussions during the preparation. We are grateful to the Referees, for carefully reading the manuscript and constructively sharing their comments on the topic with us. This work has been funded by the Kavli Foundation and the Norwegian Research Council Centre of Excellence scheme.

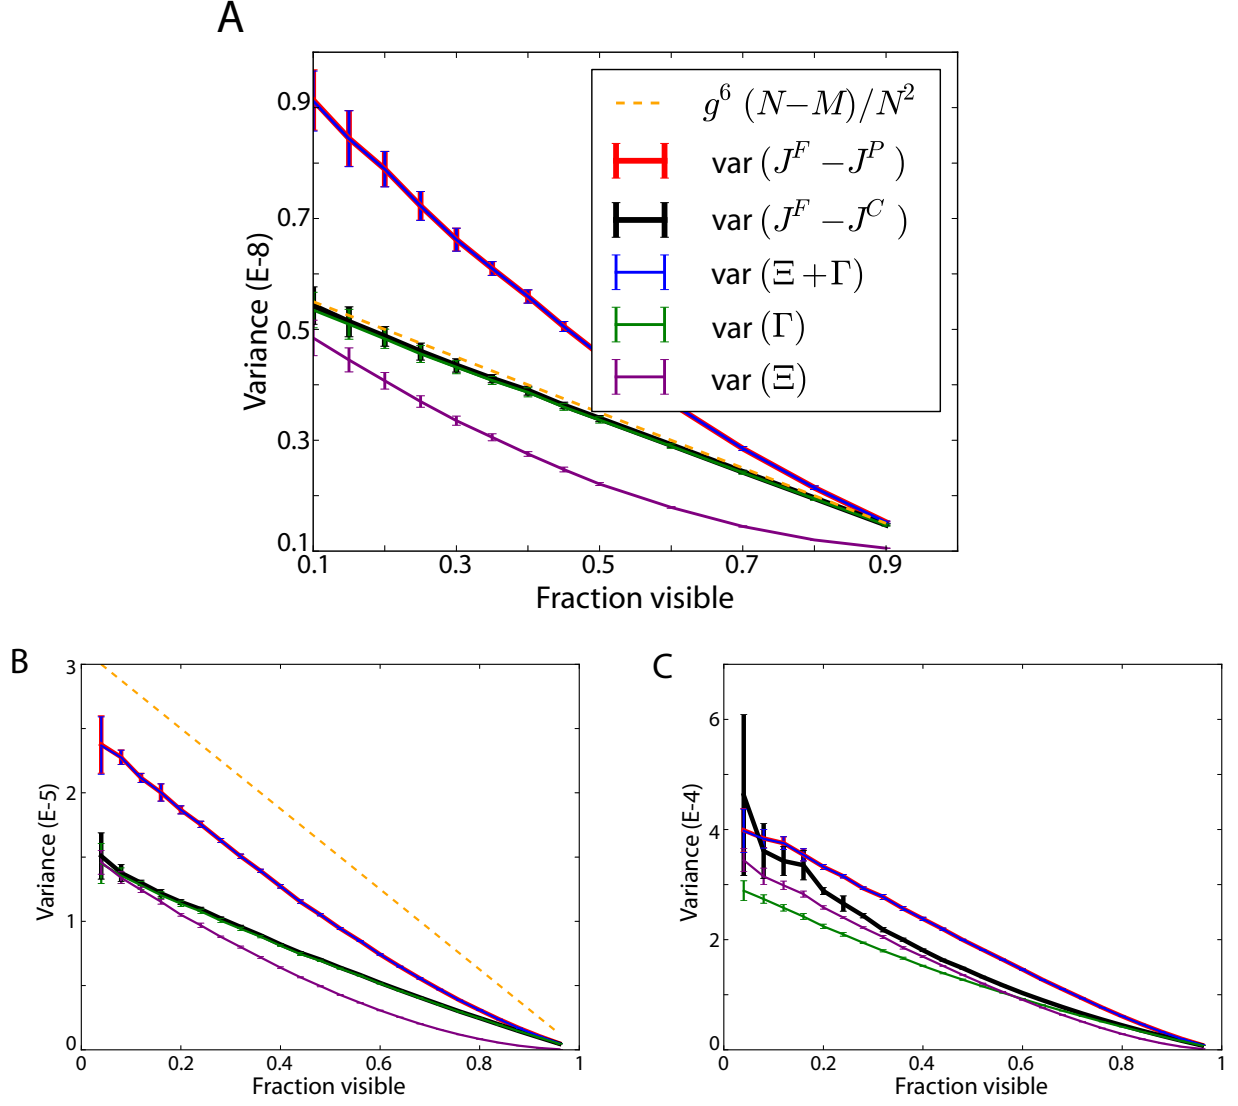


Figure 2: Correcting for the unobserved. The systematic bias, Ξ , is partially removed by solving the self-consistent equation 21. The resulting corrected couplings, J^C , provide a more accurate estimation of the ground truth with the remaining error estimated by $\text{var}(\Gamma)$, i.e. $\text{var}(J^F - J^P) \approx \text{var}(\Xi + \Gamma)$ and $\text{var}(J^F - J^C) \approx \text{var}(\Gamma)$. (A) For weak couplings, $J_0 = 0.1$, the approximations developed here perform very well, as expected. For stronger couplings $J_0 = 0.5$ (B) and $J_0 = 1$ (C) the corrected couplings do result in an improved reconstruction in every case except for the strongly coupled and heavily subsampled regime. Each inference was done 50 times, randomly selecting hidden nodes each time, and the resulting variances averaged. Error bars represent one standard deviation. The data was generated from a network with 200 nodes and a data length $L = 1.5 \times 10^{10}$ [†] for $J_0 = 0.1$ (A) and 500 neurons with $L = 10^7$ [†], for $J_0 = 0.5, 1$ (B-C).

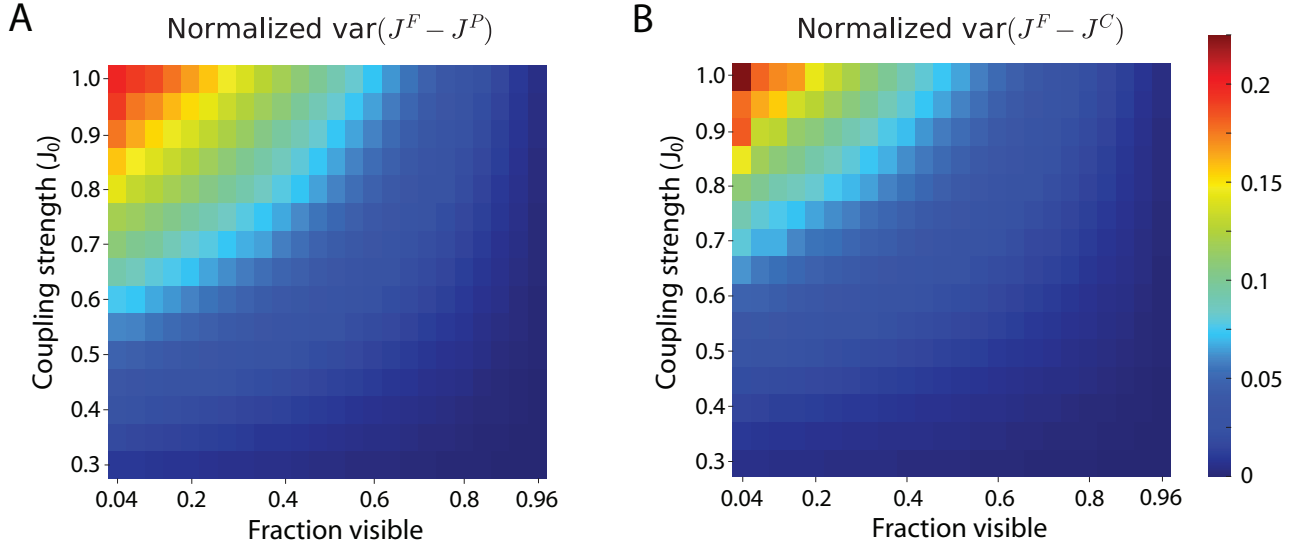


Figure 3: Normalized relative errors for the couplings inferred from the partially observed system (**A**) and using the approach described in equation (21) (**B**) as a function of subsampling (x-axis) and couplings strength J_0 (y-axis). Here the error terms are averages over 50 random selections of the hidden nodes as in Figure 2, normalized by the variance of the couplings (J_0^2/N). It is interesting to note that the simple correction is able to significantly improve the inference for all but the strongly-coupled, highly-subsampled regime.

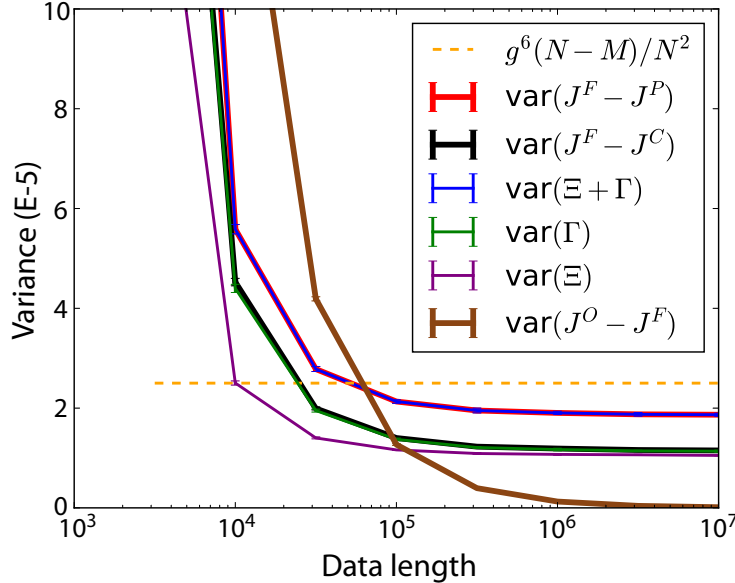


Figure 4: Hidden nodes problem as a function of data length. With increasing data length, the estimates and error terms from Figure 2 continue to decrease until saturation. It is clear that even with infinite sampling the resulting errors due to subsampling would persist while for the fully observed system (brown) the overall error tends to zero, as one would expect. It is important to note that the problem of subsampling quickly becomes relevant, with, in this example, the error due to subsampling nodes surpassing the error due to data length at below 10^5 . In this example, $J_0 = 0.5$ with a visible fraction of 0.2.

References

- [1] L. Bachschmid-Romano and M. Opper. Inferring hidden states in a random kinetic ising model: replica analysis. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(6):P06013, 2014.
- [2] C. Battistin, J. Hertz, J. Tyrcha, and Y. Roudi. Belief propagation and replicas for inference and learning in a kinetic ising model with hidden spins. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(5):P05021, 2015.
- [3] S. S. Borysov, Y. Roudi, and A. V. Balatsky. Us stock market interaction network as learned by the boltzmann machine. *The European Physical Journal B*, 88(12):1–14, 2015.
- [4] J.-P. Bouchaud. Crises and collective socio-economic phenomena: simple models and challenges. *Journal of Statistical Physics*, 151(3-4):567–606, 2013.
- [5] T. Bury. Statistical pairwise interaction model of stock market. *The European Physical Journal B*, 86(3):1–7, 2013.
- [6] A. C. Coolen, R. Kühn, and P. Sollich. *Theory of neural information processing systems*. Oxford University Press, Inc., 2005.
- [7] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38, 1977.
- [8] B. Dunn, M. Mørreaunet, and Y. Roudi. Correlations and functional connections in a population of grid cells. *PLoS Comput Biol*, 11(2):e1004052, 2015.
- [9] B. Dunn and Y. Roudi. Learning and inference in a nonequilibrium ising model with hidden nodes. *Physical Review E*, 87(2):022127, 2013.
- [10] L. Fahrmeir and H. Kaufmann. Consistency and asymptotic normality of the maximum likelihood estimator in generalized linear models. *The Annals of Statistics*, pages 342–368, 1985.
- [11] J. Hertz, Y. Roudi, A. Thorning, J. Tyrcha, E. Aurell, and H. Zeng. Inferring network connectivity using kinetic ising models. *BMC Neuroscience*, 11(Suppl 1):P51, 2010.
- [12] J. Hertz, Y. Roudi, and J. Tyrcha. Ising models for inferring network structure from spike data. *Arxiv preprint arXiv:1106.1752*, 2011.
- [13] H. Huang. Effects of hidden nodes on network structure inference. *Journal of Physics A: Mathematical and Theoretical*, 48(35):355002, 2015.
- [14] H. J. Kappen and F. d. B. Rodríguez. Efficient learning in boltzmann machines using linear response theory. *Neural Computation*, 10(5):1137–1156, 1998.
- [15] J. Kulkarni and L. Paninski. Common-input models for multiple neural spike-train data. *Network: Computation in Neural Systems*, 18(4):375–407, 2007.
- [16] T. Lezon, J. Banavar, M. Cieplak, A. Maritan, and N. Fedoroff. Using the principle of entropy maximization to infer genetic interaction networks from gene expression patterns. *Proceedings of the National Academy of Sciences*, 103(50):19033–19038, 2006.
- [17] W. Little. The existence of persistent states in the brain. *Mathematical Biosciences*, 19(1):101 – 120, 1974.
- [18] H. Mahmoudi and D. Saad. Generalized mean field approximation for parallel dynamics of the ising model. *Journal of statistical mechanics*, 2014(7), 2014.
- [19] G. McLachlan and T. Krishnan. *The EM algorithm and extensions*, volume 382. John Wiley & Sons, 2007.

- [20] M. Mézard and J. Sakellariou. Exact mean-field inference in asymmetric kinetic ising systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2011:L07001, 2011.
- [21] J. Nelder and R. Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society. Series A (General)*, pages 370–384, 1972.
- [22] J. W. Pillow and P. E. Latham. Neural characterization in partially observed populations of spiking neurons. In *Advances in Neural Information Processing Systems*, pages 1161–1168, 2007.
- [23] Y. Roudi, B. Dunn, and J. Hertz. Multi-neuronal activity and functional connectivity in cell assemblies. *Current opinion in neurobiology*, 32:38–44, 2015.
- [24] Y. Roudi and J. Hertz. Dynamical tap equations for non-equilibrium ising spin glasses. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(03), 2011.
- [25] Y. Roudi and J. Hertz. Mean field theory for nonequilibrium network reconstruction. *Physical Review Letters*, 106(4):48702, 2011.
- [26] Y. Roudi, J. Tyrcha, and J. Hertz. Ising model for neural data: model quality and approximate methods for extracting functional connectivity. *Physical Review E*, 79(5):051915, 2009.
- [27] J. Sakellariou, Y. Roudi, M. Mezard, and J. Hertz. Effect of coupling asymmetry on mean-field solutions of the direct and inverse sherrington–kirkpatrick model. *Philosophical Magazine*, 92(1-3):272–279, 2012.
- [28] E. Schneidman, M. Berry, R. Segev, and W. Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006.
- [29] D. Sherrington and S. Kirkpatrick. Solvable model of a spin-glass. *Physical review letters*, 35(26):1792, 1975.
- [30] I. H. Stevenson and K. P. Kording. How advances in neural recording affect data analysis. *Nature neuroscience*, 14(2):139–142, 2011.
- [31] R. Sundberg. Maximum likelihood theory for incomplete data from an exponential family. *Scandinavian Journal of Statistics*, pages 49–58, 1974.
- [32] T. Tanaka. Mean-field theory of boltzmann machine learning. *Physical Review E*, 58(2):2302, 1998.
- [33] J. Tyrcha and J. Hertz. Network inference with hidden units. *Mathematical biosciences and engineering: MBE*, 11(1):149–165, 2014.
- [34] A. Witoelar and Y. Roudi. Neural network reconstruction using kinetic ising models with memory. *BMC Neuroscience*, 12(Suppl 1):P274, 2011.
- [35] C. J. Wu. On the convergence properties of the em algorithm. *The Annals of statistics*, pages 95–103, 1983.
- [36] H.-L. Zeng, R. Lemoy, and M. Alava. Financial interaction networks inferred from traded volumes. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(7):P07008, 2014.