

Probabilistic inference of binary Markov random fields in spiking neural networks through mean-field approximation

Yajing Zheng^{a,b}, Shanshan Jia^{a,b}, Zhaofei Yu^{a,b,*}, Tiejun Huang^{a,b}, Jian K. Liu^c,
Yonghong Tian^{a,b,**}

^a National Engineering Laboratory for Video Technology, Department of Computer Science and Technology, Peking University, Beijing 100871, China

^b Peng Cheng Laboratory, Shenzhen 518055, China

^c Centre for Systems Neuroscience, Department of Neuroscience, Psychology and Behaviour, University of Leicester, Leicester LE1 7HA, UK

ARTICLE INFO

Article history:

Received 16 July 2019

Received in revised form 31 January 2020

Accepted 2 March 2020

Available online 9 March 2020

Keywords:

Probabilistic inference

Markov Random Fields (MRFs)

Spiking Neural Networks (SNNs)

Recurrent Neural Networks (RNNs)

Mean-field approximation

ABSTRACT

Recent studies have suggested that the cognitive process of the human brain is realized as probabilistic inference and can be further modeled by probabilistic graphical models like Markov random fields. Nevertheless, it remains unclear how probabilistic inference can be implemented by a network of spiking neurons in the brain. Previous studies have tried to relate the inference equation of binary Markov random fields to the dynamic equation of spiking neural networks through belief propagation algorithm and reparameterization, but they are valid only for Markov random fields with limited network structure. In this paper, we propose a spiking neural network model that can implement inference of arbitrary binary Markov random fields. Specifically, we design a spiking recurrent neural network and prove that its neuronal dynamics are mathematically equivalent to the inference process of Markov random fields by adopting mean-field theory. Furthermore, our mean-field approach unifies previous works. Theoretical analysis and experimental results, together with the application to image denoising, demonstrate that our proposed spiking neural network can get comparable results to that of mean-field inference.

© 2020 Elsevier Ltd. All rights reserved.

1. Introduction

The human brain is able to process information in the presence of sensory uncertainties (Meyniel, Sigman, & Mainen, 2015). For example, one can easily localize a bird in a tree via noisy visual and auditory cues. Such processes can be understood as probabilistic inference and further modeled by probabilistic graphical models (Koller & Friedman, 2009; Wainwright, Jordan, et al., 2008), including Bayesian networks and Markov Random Fields (MRFs). With an increasing volume of behavioral and physiological evidence (Doya, Ishii, Pouget, & Rao, 2007; Knill & Richards, 1996; Ma & Jazayeri, 2014; Pouget, Drugowitsch, & Kepecs, 2016) that humans do actually use probabilistic rules in perception (Kersten, Mamassian, & Yuille, 2004; Shi, Church, & Meck, 2013), sensorimotor control (Bays & Wolpert, 2007; Körding & Wolpert, 2004) and cognition (Chater, Tenenbaum, & Yuille, 2006; Jampani, Nowozin, Loper, & Gehler, 2015; Yuille & Kersten, 2006), probabilistic brain is getting recognized by

neuroscientists (Pouget, Beck, Ma, & Latham, 2013). Nevertheless, it remains unclear how the brain can perform inference. Or more precisely, how a network of spiking neurons in the brain can implement inference of probabilistic graphical models? This problem is of great importance to both computer science and brain science (Yu et al., 2020). If we know the neural algorithms of probabilistic inference, it is possible to build a machine that can perform probabilistic inference like the human brain.

In recent studies, many researchers have been devoted to developing neural circuits that can represent and implement inference of undirected probabilistic graphical models, namely MRFs (Koller & Friedman, 2009), which is widely used in computational neuroscience (Fischl et al., 2002; Ming & Hu, 2010; Probst et al., 2015; Vasta et al., 2016). The reason for focusing on MRFs is that, for directed probabilistic graphical models, one can easily convert them to MRFs via moralization (Jordan, Ghahramani, Jaakkola, & Saul, 1999; Koller & Friedman, 2009).

Here we briefly review these previous studies. Litvak and Ullman (2009) designed neural circuits to implement the operations of summation and multiplication respectively, and further implemented probabilistic computation and inference of MRFs. Steimer, Maass, and Douglas (2009) proposed using a population of spiking neurons to collect messages and another population to send messages, and then implemented the Belief Propagation (BP)

* Corresponding author at: National Engineering Laboratory for Video Technology, Department of Computer Science and Technology, Peking University, Beijing 100871, China.

** Corresponding author.

E-mail addresses: yuzf12@pku.edu.cn (Z. Yu), yhtian@pku.edu.cn (Y. Tian).




Fig. 1. A square lattice pairwise Markov random field. The filled-in circles represent the observed nodes y_i , while the empty circles represent the “hidden” nodes x_i .

algorithm, a commonly used inference method in probabilistic graphical models (Koller & Friedman, 2009; Wainwright et al., 2008). All these studies require that each neuron and synapse conduct complicated computation. However, one often observes one basic principle of the neuronal system in the brain that a single neuron or a group of neurons should work in a relatively simple style, while complex functions could be achieved when they are wired together, i.e., collaborated in a network (Buonomano & Maass, 2009; Liu & Buonomano, 2009).

In order to propose biologically more plausible neural networks to implement inference, Ott and Stoop (2007) established a relationship between the inference equation of binary MRFs and the dynamic equation of spiking neural networks through BP algorithm and reparameterization. However, their model relied on the specifically initialized messages and certain topological structures of MRFs. Yu, Chen, and Dong (2017) went a further step to relax the constraints on initialized messages, but still required the special topological structure and potential function of MRFs. Another important way is based on tree-based reparameterization algorithm (Raju & Pitkow, 2016), which, however, is only limited to the case of exponential family distributions.

In this paper, we use a mean-field approximation to treat the inference process of MRFs as a time-continuous system of a recurrent spiking neural network. We analytically prove a precise equivalence between the inference equation of Markov random fields and the dynamic equation of spiking recurrent neural networks. We show that the firing rates of neurons in the network can encode the difference between the probabilities of two states. In addition, we prove that the time course of neural firing rate can implement marginal inference of arbitrary binary Markov random fields. In this way, we can obtain the state of the neuron by counting spikes from each neuron within a time window. We further show that our mean-field approximation unifies the previous approach based on BP algorithm and reparameterization. Theoretical analysis and experimental results, together with an application to the image denoising problem, show that our proposed spiking neural network can get comparable results to that of mean-field inference.

To summarize, our contributions include the following aspects:

- We propose a spiking neural network model that can implement inference of arbitrary binary Markov random fields.

- We prove that there exists a precise equivalence between the dynamics of recurrent neural network and the inference equation of a Markov random field.
- We show that the previous approach based on BP algorithm and reparameterizations equals mean-field approximation.
- We show that our proposed spiking neural network can be used to solve practical computer vision problems, like image denoising.

The rest of the paper is organized as follows. In Section 2 we briefly review MRFs and marginal inference, then we derive the inference equation of MRFs based on mean-field approximation and show how it is related to the dynamic equation of spiking neural networks in Section 3. We show the simulation results in Section 4 and conclude in Section 5.

2. Markov random fields and marginal inference

In this section, we briefly review MRFs and marginal inference. MRFs is one typical undirected probabilistic graphical model that is widely used in computational neuroscience. Thanks to their ability to model soft contextual constraints between random variables, MRFs provide a principled probabilistic framework to model various vision problems (Chen & Tang, 2007; Dong et al., 2015; Felzenszwalb & Huttenlocher, 2006) since the visual scene modeling usually involves interactions between a subset of pixels and scene components.

In a MRF, a joint distribution $P(\{x\}) = P(x_1, x_2, \dots, x_n)$ is defined on the graph, which can be factorized into a product of potential functions according to the structure of the graph. For the MRF in Fig. 1, $P(\{x\})$ has the form:

$$P(\{x\}) = \frac{1}{Z} \prod_{(i,j) \in E} \Psi_{ij}(x_i, x_j) \prod_{i \in V} \Psi_i(x_i, y_i), \quad (1)$$

where E and V represent the set of edges and nodes in the graph respectively, $\Psi_{ij}(x_i, x_j)$ and $\Psi_i(x_i, y_i)$ denote the pairwise and unary potential functions. Z is the partition function defined as $Z = \sum_{x_1, x_2, \dots, x_n} \prod_{(i,j) \in E} \Psi_{ij}(x_i, x_j) \prod_{i \in V} \Psi_i(x_i, y_i)$. If one defines $J_{ij}(x_i, x_j) = \ln \Psi_{ij}(x_i, x_j)$ and $h_i(x_i) = \ln \Psi_i(x_i, y_i)$,¹ Eq. (1) can be rewritten as:

$$P(\{x\}) = \frac{1}{Z} \exp \left(\sum_{(i,j) \in E} J_{ij}(x_i, x_j) + \sum_{i \in V} h_i(x_i) \right) \quad (2)$$

Similar to the studies in Ott and Stoop (2007) and Yu et al. (2017), we assume that $J_{ij}(x_i, x_j) = J_{ij}x_ix_j$ and $h_i(x_i) = h_ix_i$, in which J_{ij} and h_i are constants.

The inference problems of MRFs include Maximum a Posterior (MAP) estimation and marginal inference. By MAP estimation, we refer to the estimation of a maximum of posterior point estimator. Conversely, marginal inference refers to inferring the posterior or conditional distribution over the latent causes of observations. In this paper, we only consider marginal inference. Specifically, we compute the marginal distribution of each variable x_i , that is:

$$p_i(x_i) = \sum_{x \setminus x_i} P(x_1, x_2, \dots, x_n). \quad (3)$$

¹ As the observed variable y_i is fixed, one can subsume it into the definition of $h_i(x_i)$.

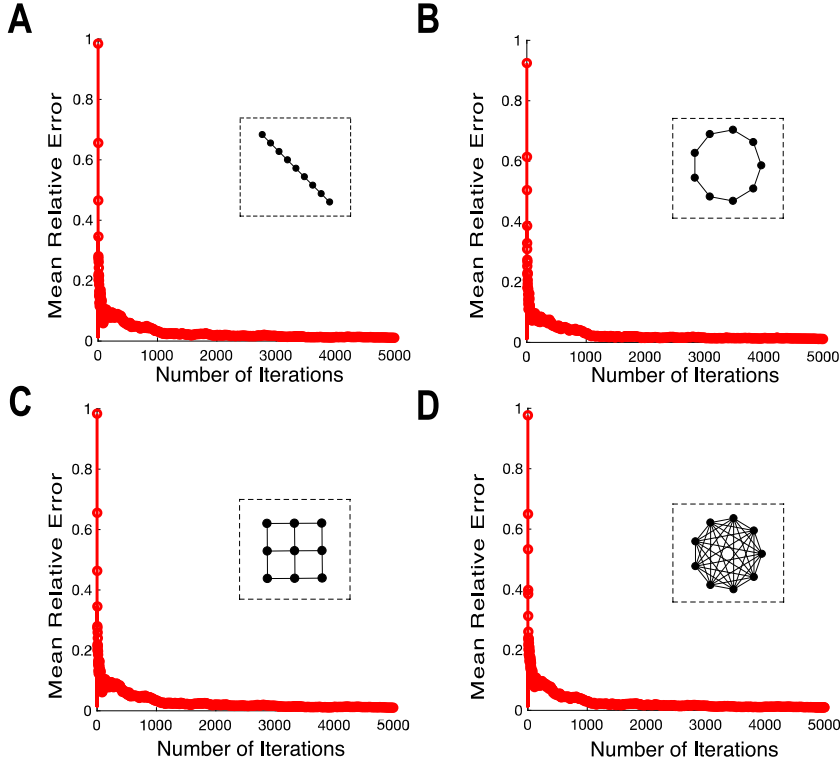


Fig. 2. Recurrent neural network achieves similar inference results as the mean-field method. The mean relative errors rapidly converge in a few iterations. All MRFs have nine nodes with different topologies as chain (A), single loop (B), grid (C) and fully connected graph (D). $\lambda_1 = \lambda_2 = 0.1$.

3. Neural implementation of marginal inference on binary MRF

In this section, we will prove that there exists a precise equivalence between the neuronal dynamics of recurrent neural networks and mean-field inference of binary MRFs. We first derive a differential equation that has the same fixed point as the mean-field inference equation of MRFs, then we show that this differential equation can be easily implemented by the dynamic equation of recurrent neural networks. In the end, we demonstrate that the previous work based on BP algorithm and reparameterization equals the mean-field approximation.

3.1. Converting mean-field inference into a differential equation

Similar to the studies in Ott and Stoop (2007) and Yu et al. (2017), we only consider inference of binary MRFs in this paper, which means the value of the variable x_i can be 1 or -1 ($x_i = 1$ or -1).

As exact inference of MRF is a NP-complete problem (Koller & Friedman, 2009), approximate inference algorithms like variational methods are often used. The main principle of variational methods is converting the inference problem to an optimization problem:

$$\min_{q(x)} KL(q(x) \parallel p(x)). \quad (4)$$

Here the target distribution $p(x)$ is approximated by a simpler distribution $q(x)$, which belongs to a family of tractable distribution. $KL(\cdot)$ represents the Kullback–Leibler divergence between two distributions. In the mean-field method, $q(x)$ is set

to be a fully factorized distribution, that is $q(x) = \prod_i b_i(x_i)$. By constraining $\sum_{x_i} b_i(x_i) = 1$ and differentiating $KL(q(x) \parallel p(x))$ with respect to $b_i(x_i)$, one can obtain the mean-field inference equation:

$$b_i^{t+1}(x_i) = \alpha \Psi_i(x_i, y_i) \exp \left(\sum_{j \in N(i)} \sum_{x_j} b_j^t(x_j) \ln \Psi_{ij}(x_i, x_j) \right), \quad (5)$$

where α is a normalization constant to make $\sum_{x_i} b_i(x_i) = 1$ and $N(i)$ denotes the set of all neighboring nodes of node i . t denotes the number of iterations, and $b_i^t(x_i)$ represents the information received by node i in the t th iteration, which is a function with respect to the state of variable x_i . When all the message converge to the fixed point, the marginal probability $p(x_i)$ can be approximated by the steady-state $b_i^\infty(x_i)$. According to the definition $\ln \Psi_{ij}(x_i, x_j) = J_{ij}(x_i, x_j) = J_{ij} x_i x_j$ and $\ln \Psi_i(x_i, y_i) = h_i(x_i) = h_i x_i$, Eq. (5) can be rewritten as:

$$b_i^{t+1}(x_i) = \alpha \exp \left(\sum_{j \in N(i)} \sum_{x_j} b_j^t(x_j) \cdot J_{ij} x_i x_j + h_i x_i \right). \quad (6)$$

In order to convert Eq. (6) to a differential equation, we reparameterize the message $b_i^t(x_i)$ of variable x_i according to:

$$n_i^t = b_i^t(x_i = 1) - b_i^t(x_i = -1), \quad (7)$$

where n_i^t can be seen as the new message received by node i in the t th iteration. Note that here the message n_i^t is independent of the state of variable x_i . When n_i^t converges to the fixed point, it can approximate the probability $p(x_i = 1) - p(x_i = -1)$. Combining Eqs. (6)–(7) and the condition $b^t(x_i = 1) +$

$b^t(x_i = -1) = 1$ defined on binary MRF, one can get that:

$$\begin{aligned}
 n_i^{t+1} &= b_i^{t+1}(x_i = 1) - b_i^{t+1}(x_i = -1) \\
 &= \alpha \exp \left(\sum_{j \in N(i)} J_{ij} \cdot (b_j^t(x_j = 1) - b_j^t(x_j = -1)) + h_i \right) \\
 &\quad - \alpha \exp \left(\sum_{j \in N(i)} -J_{ij} \cdot (b_j^t(x_j = 1) - b_j^t(x_j = -1)) - h_i \right) \\
 &= \tanh \left(\sum_{j \in N(i)} J_{ij} \cdot (b_j^t(x_j = 1) - b_j^t(x_j = -1)) + h_i \right) \\
 &= \tanh \left(\sum_{j \in N(i)} J_{ij} \cdot n_j^t + h_i \right).
 \end{aligned} \tag{8}$$

Note that the third equality of Eq. (8) holds as

$$\begin{aligned}
 \frac{1}{\alpha} &= \sum_{x_i} \exp \left(\sum_{j \in N(i)} \sum_{x_j} b_j^t(x_j) \cdot J_{ij} x_i x_j + h_i x_i \right) \\
 &= \exp \left(\sum_{j \in N(i)} J_{ij} \cdot (b_j^t(x_j = 1) - b_j^t(x_j = -1)) + h_i \right) \\
 &\quad + \exp \left(\sum_{j \in N(i)} -J_{ij} \cdot (b_j^t(x_j = 1) - b_j^t(x_j = -1)) - h_i \right).
 \end{aligned} \tag{9}$$

It is easy to prove that the following differential equation has the same fixed point as Eq. (8).

$$\tau_0 \frac{dn_i(t)}{dt} = -n_i(t) + \tanh \left(\sum_{j \in N(i)} J_{ij} \cdot n_j(t) + h_i \right), \tag{10}$$

where τ_0 is a time constant that determines the time needed for the network to reach the fixed point.

3.2. Dynamic equation of spiking recurrent neural networks

Recurrent neural networks are composed of a population of interconnected neurons, which have been widely used to model cortical response properties in computational neuroscience (Rao, 2004, 2005). Here, we drive the firing-rate based equation of spiking recurrent neural network based on two steps (Dayan & Abbott, 2001): (1) Determining how the total synaptic input to a neuron depends on the firing rate of its presynaptic afferents. (2) Modeling how the firing rate of the postsynaptic neuron depends on its total synaptic input.

First of all, considering the recurrent neural network consists of N spiking neurons z_1, z_2, \dots, z_N , the input current to the neuron z_i at time t is $I_i(t)$, which includes the recurrent input of spike sequence from other neurons and can be computed as:

$$I_i(t) = \sum_{j=1}^N w_{ij} \int_{-\infty}^t \kappa(t - \tau) S_j(\tau) d\tau, \tag{11}$$

where $S_j(\tau)$ denotes the firing spike sequence of neuron z_j defined as a sum of Dirac δ function $S_j(t) = \sum_f \delta(t - t_j^f)$, t_j^f is firing time of the f th spike of neuron z_j . w_{ij} denotes the synaptic weight between neuron z_i and z_j , $\kappa(t)$ is the synaptic kernel that describes the time course of the synaptic current in response to a presynaptic spike arriving at time t . The most frequently used form of synaptic kernel is an exponential kernel, that is, $\kappa(t) = \frac{1}{\tau_s} \exp\left(-\frac{t}{\tau_s}\right)$ with the membrane time constant τ_s .

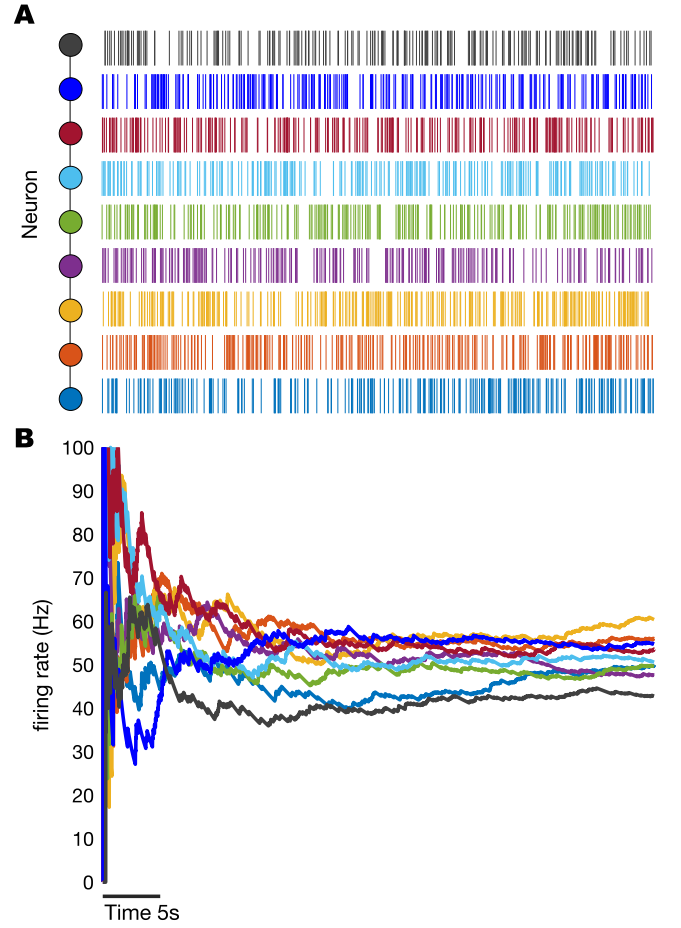


Fig. 3. Inference of a chain-structured MRF with spiking recurrent neural network. (A) Spiking activity of the neurons in a recurrent neural network. (B) Time course of firing rates of 9 neurons shown in (A).

In fact, the neural response function $S_j(t)$ could be replaced by the firing rate $r_j(t)$ of neuron z_j as $r_j(t) = \frac{1}{\Delta t} \int_t^{t+\Delta t} \langle S_j(\tau) \rangle d\tau$ with $\langle S_j(t) \rangle$ denoting the trial-average neural response function, thus Eq. (11) can be rewritten as:

$$I_i(t) = \sum_{j=1}^N w_{ij} \int_{-\infty}^t \frac{1}{\tau_s} \exp\left(-\frac{t-\tau}{\tau_s}\right) r_j(\tau) d\tau. \tag{12}$$

By taking the derivative of $I_k(t)$ with respect to time t , one obtain:

$$\tau_s \frac{dI_i(t)}{dt} = -I_i(t) + \sum_{j=1}^N w_{ij} r_j(t), \tag{13}$$

with τ_s denoting the time constant that describes the decay of the synaptic conductance.

So far we can determine the input current to postsynaptic neuron in terms of the firing rates of the presynaptic neurons. To obtain the firing-rate model, we also need to determine the postsynaptic firing rate with the current $I_i(t)$. For time-independent inputs, the firing rate $r_i(t)$ of the postsynaptic neuron z_i can be expressed as $r_i(t) = F(I_i(t))$, where $F(x)$ denotes the neuronal activation function. As the firing rate does not follow changes of the total synaptic current instantaneously, the firing rate is often modeled by a low-pass filtered version of the synaptic current:

$$\tau_r \frac{dr_i(t)}{dt} = -r_i(t) + F(I_i(t)). \tag{14}$$

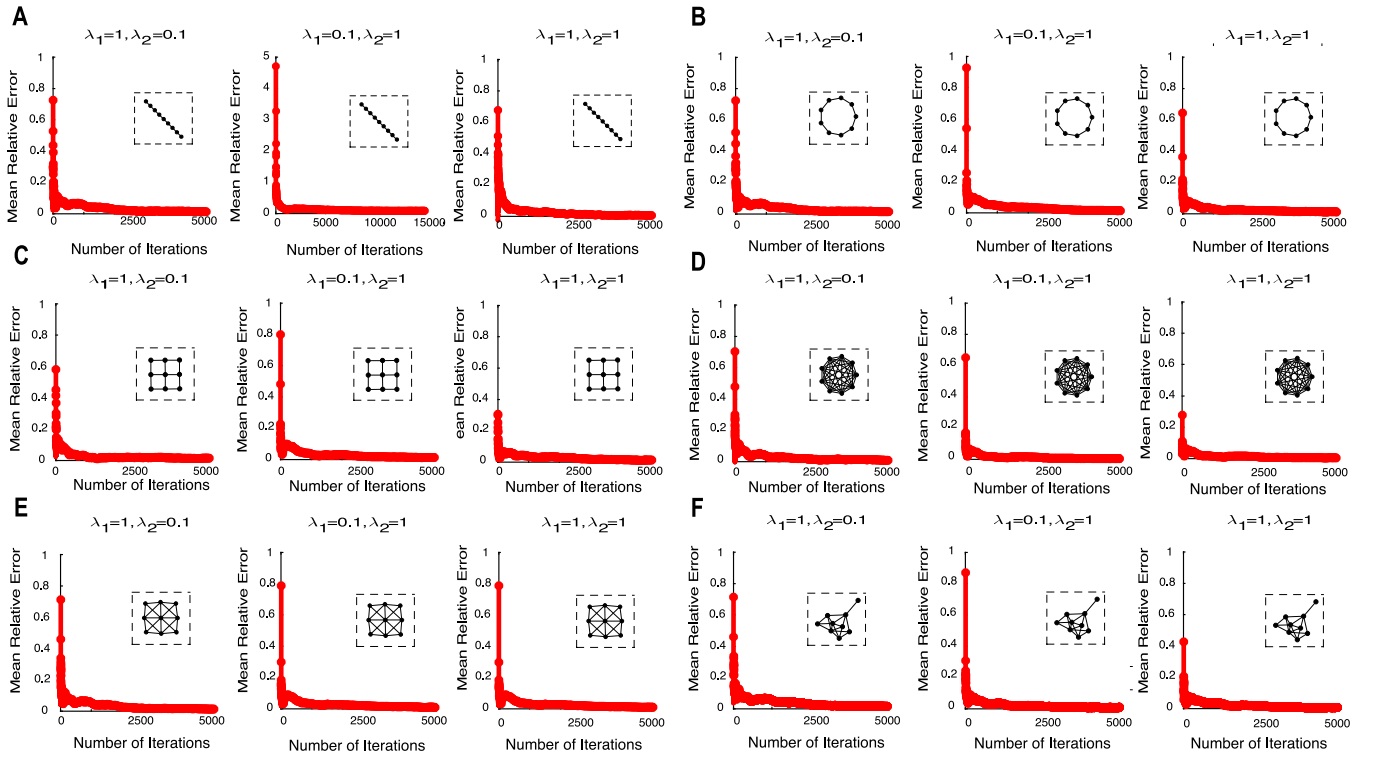


Fig. 4. Illustration of the robustness of our method with different settings of parameters on different topologies as chain (A), single loop (B), grid (C), fully connected graph (D), 8-connected grid topology (E) and random connected topology with connection probability 0.5 for each edge (F).

Under the constraints of time-independent inputs, the steady state of the postsynaptic current $I_i(t)$ is $\lim_{t \rightarrow \infty} I_i(t) = \sum_{j=1}^N w_{ij} r_j(t)$. If $\tau_r \gg \tau_s$, we can make the approximation that Eq. (13) comes to equilibrium quickly compared to Eq. (14). Consequently, we can further replace $I_i(t)$ by $\sum_{j=1}^N w_{ij} r_j(t)$ in Eq. (14) and obtain:

$$\tau_r \frac{dr_i(t)}{dt} = -r_i(t) + F \left(\sum_{j=1}^N w_{ij} r_j(t) \right). \quad (15)$$

In recurrent neural networks, except for the input current from recurrent neurons, there also exists an external input current. Incorporating the external input current $I_i^{ext}(t)$ to Eq. (15), the firing rate of the recurrent neuron k is determined by:

$$\tau_r \frac{dr_i(t)}{dt} = -r_i(t) + F \left(I_i^{ext}(t) + \sum_{j=1}^N w_{ij} r_j(t) \right). \quad (16)$$

3.3. Implementation of inference with neural network

Now one can relate the inference equation of MRFs (Eq. (10)) to the dynamics of recurrent neural networks (Eq. (16)). Obviously, Eq. (10) is equivalent to Eq. (16) if the following equations hold:

$$\tau_r = \tau_0, \quad (17)$$

$$r_i(t) = n_i(t), \quad (18)$$

$$w_{ij} = \begin{cases} J_{ij} & \text{if } j \in N(i) \\ 0 & \text{others} \end{cases}, \quad (19)$$

$$I_i^{ext}(t) = h_i, \quad (20)$$

$$F(x) = \tanh(x). \quad (21)$$

Eqs. (17)–(21) mean that if the synaptic weights w_{ij} and input current $I_i^{ext}(t)$ of a recurrent neural network encode the potential

functions J_{ij} and h_i of a binary MRF respectively, the firing rate $r_i(t)$ of neuron z_i encodes the probability $p(x_i = 1) - p(x_i = -1)$. Moreover, the time course of neural firing rate in the recurrent neural network can implement marginal inference of the MRF. Thus, we can read out the inference result by counting spikes from each neuron within a time window. Note that as the value of $n_i(t)$ varies from -1 to 1 , the firing rate $r_i(t)$ in Eq. (18) could be negative, which is biologically implausible. As discussed in Rao (2004), we can assume that the actual firing rate $\hat{r}_i(t)$ is linearly related to the “firing rate” $r_i(t)$ obtained from Eq. (16), that is, $\hat{r}_i(t) = ar_i(t) + b$. Here a is a positive factor and b is a rectification value that ensure $\hat{r}_i(t)$ to be positive. In conclusion, we implement mean-field inference of binary MRFs with spiking recurrent neural networks.

3.4. Relating mean-field inference to belief propagation

Here we will build the relationship between mean-field inference and BP, and show that the previous work based on BP and reparameterization equals the mean-field inference.

Previous studies have tried to relate BP algorithm of binary MRF to the dynamics of Hopfield Networks by deriving a new formulation of belief propagation based on reparameterization (Ott & Stoop, 2007; Yu et al., 2017):

$$\mu_i^{t+1} = \tanh \left(\sum_{j \in N(i)} \tanh^{-1} \left[\tanh(J_{ij}) \tanh \left(\sum_{s \in N(j) \setminus i} n_{s \rightarrow j}^t + h_j \right) \right] + h_i \right), \quad (22)$$

where μ_i^{t+1} represents the new message after reparameterization of node i at the $t + 1$ th iteration, and $\mu_i^\infty = p(x_i = 1) - p(x_i = -1)$. $n_{s \rightarrow j}^t$ is a function of the message $m_{s \rightarrow j}^t(x_j)$ in BP

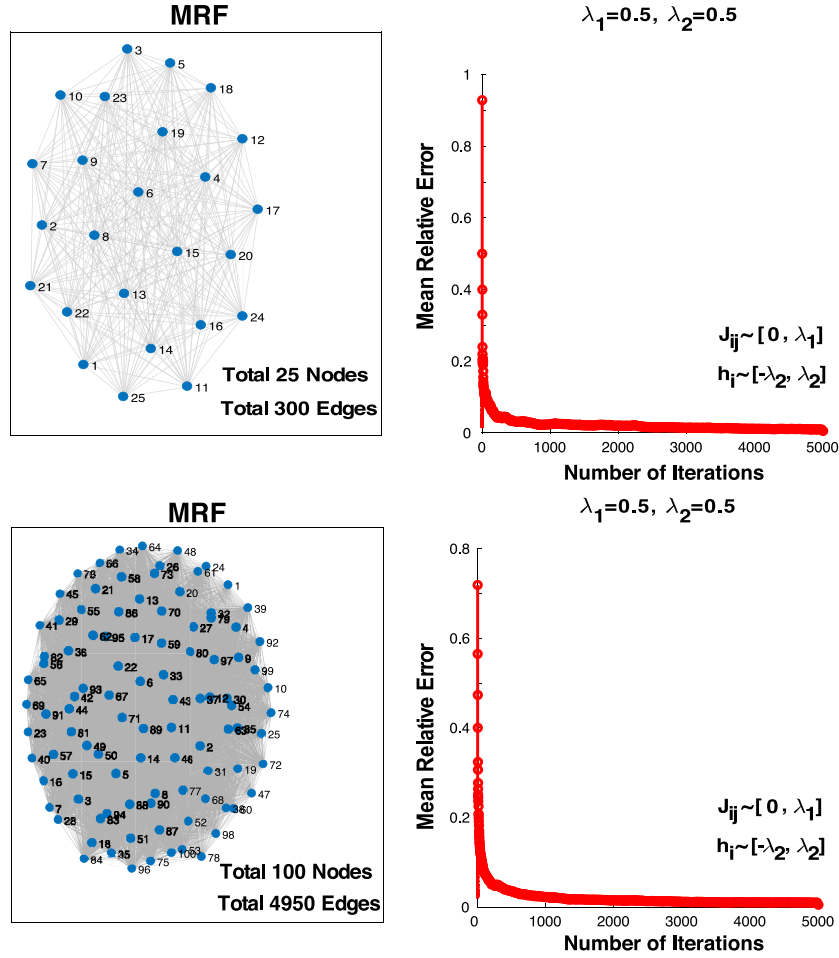


Fig. 5. Inference performance of recurrent neural networks on large-scale MRFs.

that is sent from node s to node j in the t th iteration. To be specific, $n_{s \rightarrow j}^t = \tanh^{-1}(m_{s \rightarrow j}^t(x_j = 1) - m_{s \rightarrow j}^t(x_j = -1))$. With the assumptions that the number of neighboring nodes of each node is large enough ($N(j) \gg 1$) and the potential function is small ($J_{ij} \ll 1$ and $h_i \ll 1$), Ott and Stoop (2007) and Yu et al. (2017) proved that Eq. (22) can be simplified to:

$$\mu_i^{t+1} = \tanh\left(\sum_{j \in N(i)} \tanh(J_{ij}) \cdot \mu_j^t + h_i\right). \quad (23)$$

As $J_{ij} \ll 1$, $\tanh(J_{ij}) \approx J_{ij}$. Thus one can further simplify Eq. (23) to:

$$\mu_i^{t+1} = \tanh\left(\sum_{j \in N(i)} J_{ij} \cdot \mu_j^t + h_i\right). \quad (24)$$

One can find that there exists a precise equivalence between Eqs. (24) and (8), which implies that the previous work based on BP and reparameterization equals the mean-field approximation. These results suggest that the Hopfield networks used in the previous work actually implement mean-field inference, instead of the BP algorithm. In addition, our current results explain the experiments in Yu et al. (2017) where the inference result based on Hopfield networks is not as accurate as that of BP when the potential function is large ($J_{ij} > 1$ and $h_i > 1$). These errors come from the difference between mean-field inference and the BP algorithm.

4. Simulation experiments

To validate the proposed computational framework, we evaluate the performance of recurrent neural networks through simulation experiments. We firstly test the accuracy of the propose method, and then prove that it is robust to different parameters. At last, we scale up the proposed spiking neural network to solve practical computer vision problems.

4.1. Testing on the accuracy of our method

In order to test the accuracy of the proposed method, we generated several MRFs with different graph topologies (chain, single loop, grid and fully connected graph, see Fig. 2), and perform inference of these MRFs with spiking recurrent neural network and mean-field method respectively.

For a MRF with M nodes, we calculated marginal probabilities for all these M nodes with mean-field method and the corresponding recurrent neural network respectively. The mean relative error δ is defined as follows:

$$\delta = \frac{1}{M} \sum_{i=1}^M \frac{|P^{MF}(x_i = 1) - P^{RNN}(x_i = 1)|}{P^{MF}(x_i = 1)}, \quad (25)$$

where $P^{MF}(x_i = 1)$ represents the marginal probabilities computed with mean-field method, and $P^{RNN}(x_i = 1)$ represents the result obtained by the corresponding recurrent neural network.

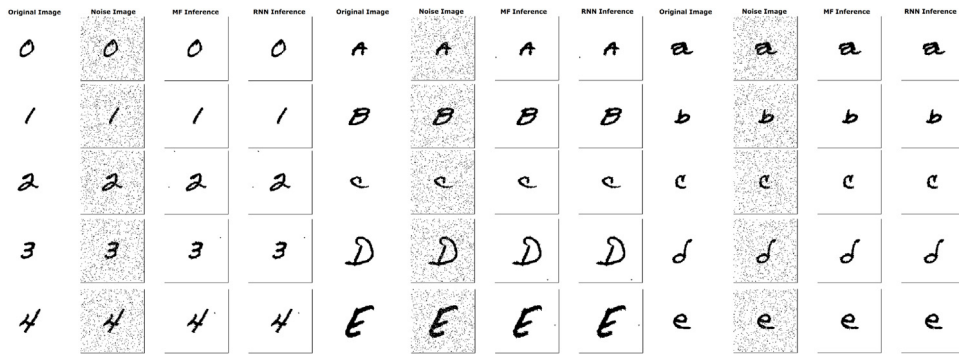


Fig. 6. Some examples of image denoising on NIST SD 19 dataset with mean-field inference and recurrent neural networks.

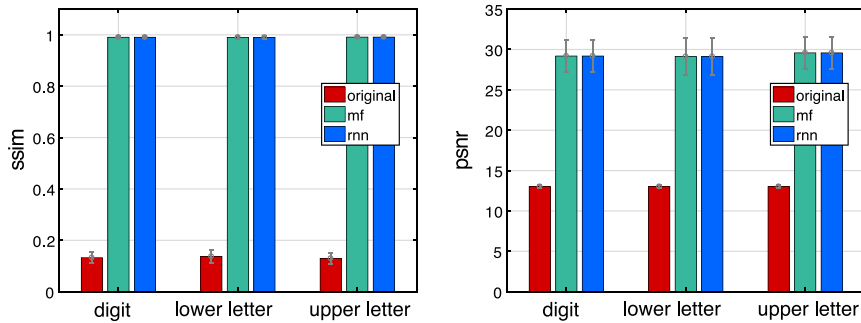


Fig. 7. Comparison of denoising result on NIST SD 19 Database with different methods. There are 62 categories from the dataset (10 classes for the digit, 26 kinds of upper letters and 26 lower letters). Average results over 5 independent trials are shown.

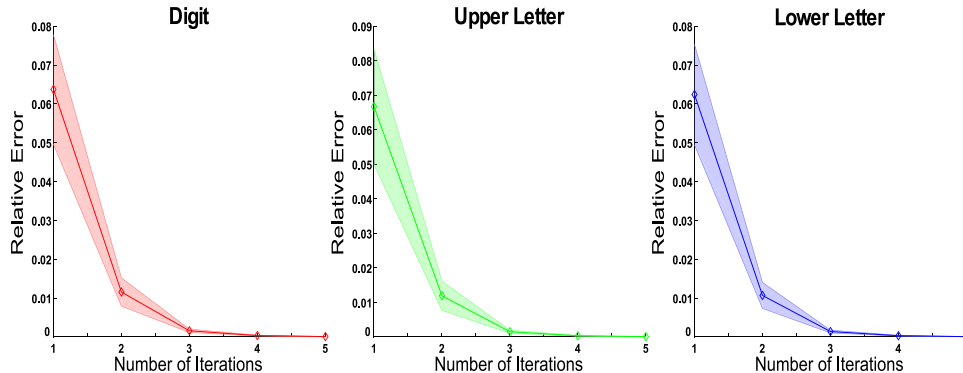


Fig. 8. Recurrent neural networks achieve similar results as those of mean-field inference. The relative error decays to 0 rapidly. The red, green and blue curves represent the results for digit, upper letter and lower letter respectively. All the results are averaged over 5 independent trials, shaded area indicates standard deviation (STD). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Fig. 2 illustrates how the relative errors rapidly convergence with a few iterations. For each MRF, the potential functions J_{ij} and h_i are drawn from two uniform distributions on $[0, \lambda_1]$ and $[-\lambda_2, \lambda_2]$ respectively. One can find that even for MRFs with different topologies, the error decreases in a fast way with only a few iterations. These results imply that the simulation of spiking recurrent neural networks can get comparable results as an analytical mean-field method.

To illustrate the inference mechanism of the spiking recurrent neural network, Fig. 3A shows the spiking activity of all 9 neurons in the recurrent neural network when performing inference of a 9-node MRF with chain structure (Fig. 2A). Here the mapping between actual firing rate $\hat{r}_i(t)$ and the “firing rate” $r_i(t)$ is $\hat{r}_i(t) = 50r_i(t) + 50$. Thus the maximum firing rate of each neuron is 100 Hz. Fig. 3B shows the time course of the firing rate of each neuron. One can see that the firing rate of each neuron converges to a fixed value and then fluctuates around it.

4.2. Testing on the robustness of our method

The experimental results above indicate that the inference model of recurrent neural networks can get accurate results as mean-field inference for a given set of parameters of λ_1 and λ_2 as 0.1. Here we make a concrete analysis of the robustness of our model with different parameters. Fig. 4 shows the results where λ_1 and λ_2 are set to different combinations of 1 and 0.1, except the setting that $\lambda_1 = \lambda_2 = 0.1$ as shown in Fig. 2. We can see that, in all cases, the errors converge to almost zero in a fast manner. These results indicate that, different from the previous works (Ott & Stoop, 2007; Yu et al., 2017) that only apply to MRFs with special potential function ($J_{ij} \ll 1$ and $h_i \ll 1$), our method is robust to different parameters and could implement inference for arbitrary MRFs.

Then we investigate whether our framework can be scaled up to large-scale MRFs with more nodes. Two examples are included

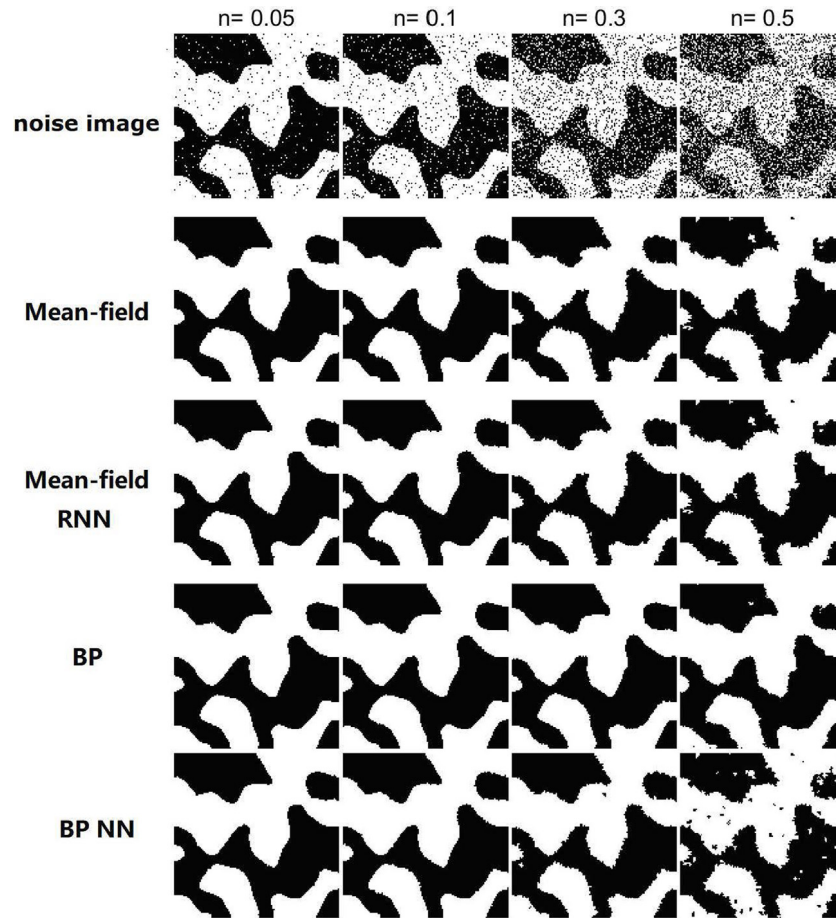


Fig. 9. Image denoising with mean-field inference, recurrent neural networks, BP algorithm and BP-based neural networks. Here n denotes different noise levels.

here: a MRF with 25 nodes and 300 edges and a MRF with 100 nodes and 4950 edges. As shown in Fig. 5, the same conclusion is obtained that the spiking recurrent neural networks can get comparable results as the mean-field method.

4.3. Binary images denoising by recurrent neural networks

Here we investigate whether our spiking neural network can be scaled up to solve more realistic tasks. We consider the task of image denoising, that is, correcting an image that has been corrupted. In the field of image processing, the researchers often model image denoising problem by MRFs with grid-like structures (shown in Fig. 1) and then convert the denoising problem to MAP estimation or marginal inference problem. Based on this, we can also tackle this problem with recurrent neural networks by computing the marginal probabilities of each pixel and then infer whether this pixel is white or black in a binary setting.

The image denoising experiments are performed on the NIST Special Database 19 (SD 19), which contains NIST's entire corpus of training materials for handwritten document and character recognition. This dataset includes samples from 3600 writers, consisting of 10 digits 0–9, 26 lower letters a–z and 26 upper letters A–Z. Therefore we have totally 62 categories. During the experiment, 100 images of each class are randomly selected as dataset. All images used here are 128×128 pixels. In this experiment, each image is modeled by a square lattice pairwise MRF (shown in Fig. 1), where the hidden variables $\{x\} = \{x_1, x_2, \dots, x_n\}$ represent the denoise image and observed variables $\{y\} = \{y_1, y_2, \dots, y_n\}$ represent the observed noise image. As observed pixel value is usually the same as the true pixel value,

so the unary potential $h(x_i)$ is set to 0.1 if the variable x_i is the same as the observation variable y_i and -0.1 otherwise, that is, $h_i = 0.1$. Besides, as nearby pixel values are usually the same in an image, the pairwise potential function $J_{ij}(x_i, x_j)$ is set to 0.8 if $x_i = x_j$ and -0.8 otherwise ($J_{ij} = 0.8$). All the other settings in this experiment are the same as in experiment 4.1 above.

Fig. 6 shows some examples of image denoising with mean-field inference and the corresponding recurrent neural network. Here the noise images are generated by randomly flipping the pixel value with a probability of 5%.

We also quantitatively analyze these results by computing the structural similarity index (SSIM) and the peak signal-to-noise ratio (PSNR). As shown in Fig. 7, the SSIM of the original image, denoised image by mean-field inference, and denoised image by recurrent neural networks are 13.01 ± 0.14 , 29.19 ± 1.97 and 29.19 ± 1.97 , respectively. The PSNR of the original image, denoised image by mean-field inference, and denoised image by recurrent neural networks are 0.1322 ± 0.0207 , 0.9905 ± 0.0047 and 0.9905 ± 0.0047 respectively. All these results demonstrate that recurrent neural networks can get the same denoising results as mean-field inference. Fig. 8 illustrates how the mean relative error between recurrent neural networks and mean-field inference varies over time. We can find the error converges to 0 with a few iterations.

4.4. Comparison among different neural network based image denoising methods

In Section 3.4, we have proved that the previous approaches based on BP and reparameterization (BP-based neural networks)




Fig. 10. Comparison of the performance of difference denoising methods with respect to noise level. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

can be unified in our framework. In order to test this, we compare our method with the BP algorithm and the BP-based neural network model for the task of image denoising. In order to increase the difficulty of inference, here we created a dataset of 100 images with 128×128 pixels by making randomly noisy images and then smooth them to get true output values. Fig. 9 shows one example of the randomly generated binary images. One can find that there exists more separated space in these images compared with the images in NIST SD 19. Thus it is more difficult to be denoised.

To compare the performance of these algorithms, we add the different levels of *salt and pepper* noise on the binary images, and characterize the quality of the denoised images with the criterion of the SSIM and PSNR. Fig. 9 illustrates one example of image denoising with mean-field inference, recurrent neural networks, BP algorithm (Koller & Friedman, 2009), and BP-based neural networks (Yu et al., 2017). Fig. 10 compares the SSIM and PSNR of different methods and noise levels. One can see that the performance of recurrent neural networks (purple curve) is the same as mean-field inference (red curve). Besides, one can also find that the performance of BP-based neural networks (green curve) is nearly the same as that of recurrent neural networks (purple curve) and mean-field inference (red curve), which demonstrates that the previous work equals mean-field inference and can be unified in our framework. Note that there exists a gap between BP-based neural networks (green curve) and recurrent neural networks (purple curve) when the noise level is larger than 0.3, which comes from the approximation between Eqs. (23) and (24).

5. Conclusion

In this paper, we prove that there exists a precise equivalence between the dynamics of recurrent neural network and mean-field inference of binary Markov random fields. We show that if the synaptic weights and input current encode the potential function of MRFs, the firing rates of neuron in recurrent neural networks encode the difference between the probabilities for two states. The time course of neuronal firing rate can implement marginal inference. Theoretical analysis and experiments on MRFs with different topologies show that our neural network can get the same performance as the mean-field method. Besides, we also apply our proposed spiking framework to practical computer vision problem, i.e., binary images denoising.

Differ from previous works based on BP algorithm and reparameterization, where the potential functions of MRF should meet some strict conditions, we design a spiking network that can

implement mean-field inference for arbitrary MRFs. What's more, we have demonstrated that our work unifies previous works.

The previous work of neural implementation of Bayesian inference (Deneve, 2008; Rao, 2004; Yu, Chen, & Deng, 2018) with recurrent neural networks focused on inference of hidden Markov models. There also exist some studies (Guo, Yu, Deng, Hu, & Chen, 2017; Rao, 2005) that extended the networks to a multilayer structure to perform hierarchical Bayesian inference. Different from these works, we are focusing on how spiking neural networks are able to implement probabilistic inference of MRF. In future work, we will try to extend our proposed framework to tackle more advanced realistic problems, like recognition and stereo matching.

Acknowledgments

This work is supported in part by the National Natural Science Foundation of China under grants 61806011, 61825101 and U1611461, in part by National Postdoctoral Program for Innovative Talents, China under grant BX20180005, in part by China Postdoctoral Science Foundation under grant 2018M630036, in part by the Zhejiang Lab, China under grants 2019KC0AB03 and 2019KC0AD02, in part by the Royal Society Newton Advanced Fellowship under grant NAF-R1-191082.

References

- Bays, Paul M., & Wolpert, Daniel M. (2007). Computational principles of sensorimotor control that minimize uncertainty and variability. *The Journal of Physiology*, 578(2), 387–396.
- Buonomano, Dean V., & Maass, Wolfgang (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, 10(2), 113.
- Chater, Nick, Tenenbaum, Joshua B., & Yuille, Alan (2006). Probabilistic models of cognition: where next?. *Trends in Cognitive Sciences*, 10(7), 292–293.
- Chen, Jia, & Tang, Chi-Keung (2007). Spatio-temporal Markov Random Field for video denoising. In *IEEE conference on computer vision and pattern recognition* (pp. 1–8). IEEE.
- Dayan, Peter, & Abbott, Laurence F. (2001). *Theoretical neuroscience*, Vol. 806. Cambridge, MA: MIT Press.
- Deneve, Sophie (2008). Bayesian spiking neurons i: inference. *Neural Computation*, 20(1), 91–117.
- Dong, Jianwu, Liu, Tian, Chen, Feng, Zhou, Dong, Dimov, Alexey, Raj, Ashish, et al. (2015). Simultaneous phase unwrapping and removal of chemical shift (SPURS) using graph cuts: application in quantitative susceptibility mapping. *IEEE Transactions on Medical Imaging*, 34(2), 531–540.
- Doya, Kenji, Ishii, Shin, Pouget, Alexandre, & Rao, Rajesh P. N. (2007). *Bayesian brain: Probabilistic approaches to neural coding*. MIT press.
- Felzenszwalb, Pedro F., & Huttenlocher, Daniel P. (2006). Efficient belief propagation for early vision. *International Journal of Computer Vision*, 70(1), 41–54.

- Fischl, Bruce, Salat, David H, Busa, Evelina, Albert, Marilyn, Dieterich, Megan, Haselgrove, Christian, Van Der Kouwe, Andre, Killiany, Ron, Kennedy, David, Klaveness, Shuna, et al. (2002). Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33(3), 341–355.
- Guo, Shangqi, Yu, Zhaofei, Deng, Fei, Hu, Xiaolin, & Chen, Feng (2017). Hierarchical bayesian inference and learning in spiking neural networks. *IEEE Transactions on Cybernetics*, 49(1), 133–145.
- Jampani, Varun, Nowozin, Sebastian, Loper, Matthew, & Gehler, Peter V (2015). The informed sampler: A discriminative approach to bayesian inference in generative computer vision models. *Computer Vision and Image Understanding*, 136, 32–44.
- Jordan, Michael I, Ghahramani, Zoubin, Jaakkola, Tommi S, & Saul, Lawrence K (1999). An introduction to variational methods for graphical models. *Machine Learning*, 37(2), 183–233.
- Kersten, Daniel, Mamassian, Pascal, & Yuille, Alan (2004). Object perception as bayesian inference. *Annual Review of Psychology*, 55, 271–304.
- Knill, David C., & Richards, Whitman (1996). *Perception as Bayesian inference*. Cambridge University Press.
- Koller, Daphne, & Friedman, Nir (2009). *Probabilistic graphical models: principles and techniques*. MIT press.
- Körding, Konrad P., & Wolpert, Daniel M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244.
- Litvak, Shai, & Ullman, Shimon (2009). Cortical circuitry implementing graphical models. *Neural Computation*, 21(11), 3010–3056.
- Liu, J. K., & Buonomano, D. V. (2009). Embedding multiple trajectories in simulated recurrent neural networks in a self-organizing manner. *Journal of Neuroscience*, 29(42), 13172–13181.
- Ma, Wei Ji, & Jazayeri, Mehrdad (2014). Neural coding of uncertainty and probability. *Annual Review of Neuroscience*, 37, 205–220.
- Meyniel, Florent, Sigman, Mariano, & Mainen, Zachary F. (2015). Confidence as Bayesian probability: from neural origins to behavior. *Neuron*, 88(1), 78–92.
- Ming, Yansheng, & Hu, Zhanyi (2010). Modeling stereopsis via Markov random field. *Neural Computation*, 22(8), 2161–2191.
- Ott, Thomas, & Stoop, Ruedi (2007). The neurodynamics of belief propagation on binary Markov Random Fields. In *Advances in neural information processing systems* (pp. 1057–1064).
- Pouget, Alexandre, Beck, Jeffrey M, Ma, Wei Ji, & Latham, Peter E (2013). Probabilistic brains: knowns and unknowns. *Nature Neuroscience*, 16(9), 1170.
- Pouget, Alexandre, Drugowitsch, Jan, & Kepecs, Adam (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nature Neuroscience*, 19(3), 366–374.
- Probst, Dimitri, Petrovici, Mihai A, Bytschok, Ilja, Bill, Johannes, Pecevski, Dejan, Schemmel, Johannes, et al. (2015). Probabilistic inference in discrete spaces can be implemented into networks of LIF neurons. *Frontiers in Computational Neuroscience*, 9, 13.
- Raju, Rajkumar Vasudeva, & Pitkow, Xaq (2016). Inference by reparameterization in neural population codes. In *Advances in neural information processing systems* (pp. 2029–2037).
- Rao, Rajesh P. N. (2004). Bayesian computation in recurrent neural circuits. *Neural Computation*, 16(1), 1–38.
- Rao, Rajesh P. (2005). Hierarchical bayesian inference in networks of spiking neurons. In *Advances in neural information processing systems* (pp. 1113–1120).
- Shi, Zhuanghua, Church, Russell M., & Meck, Warren H. (2013). Bayesian optimization of time perception. *Trends in Cognitive Sciences*, 17(11), 556–564.
- Steimer, Andreas, Maass, Wolfgang, & Douglas, Rodney (2009). Belief propagation in networks of spiking neurons. *Neural Computation*, 21(9), 2502–2523.
- Vasta, Roberta, Augimeri, Antonio, Cerasa, Antonio, Nigro, Salvatore, Gramigna, Vera, Nonnis, Matteo, et al. (2016). Hippocampal subfield atrophies in converted and not-converted mild cognitive impairments patients by a Markov random fields algorithm. *Current Alzheimer Research*, 13(5), 566–574.
- Wainwright, Martin J., Jordan, Michael I., et al. (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2), 1–305.
- Yu, Zhaofei, Chen, Feng, & Deng, Fei (2018). Unification of MAP estimation and marginal inference in recurrent neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 29(11), 5761–5766.
- Yu, Zhaofei, Chen, Feng, & Dong, Jianwu (2017). Neural network implementation of inference on binary Markov Random Fields with probability coding. *Applied Mathematics and Computation*, 301, 193–200.
- Yu, Zhaofei, Liu, Jian K., Jia, Shanshan, Zhang, Yichen, Zheng, Yajing, Tian, Yonghong, et al. (2020). Towards the next generation of retinal neuroprosthesis: Visual computation with spikes. arXiv preprint arXiv:2001.04064.
- Yuille, Alan, & Kersten, Daniel (2006). Vision as Bayesian inference: analysis by synthesis?. *Trends in Cognitive Sciences*, 10(7), 301–308.