

Antagonistic cross-links and reciprocal couplings emerge in optimal multisensory integration

He Wang¹, Wen-Hao Zhang², K. Y. Michael Wong¹, and Si Wu³

¹Hong Kong University of Science and Technology, ²University of Pittsburgh, ³Peking University.

Abstract

Information from multiple sources is exploited by the neural system to facilitate reliable and flexible information processing, in a near-optimal way as predicted by Bayes' rule. However, the underlying network architecture achieving the optimal integration is largely unknown, especially when the prior distribution describing the multiple sensory stimuli is only partially correlated. Here, we study a decentralized architecture, in which each module is a recurrent neural network processing input from one source, and cross-talks among them can facilitate integration. To achieve inter-modular communication, each module receives direct input from other sources through the *cross-links*, and indirect input from the other modules through *reciprocal links*. Through theoretical analysis and network optimization, we investigate how multisensory likelihoods and priors are encoded in different components of the network structure. We found that the multisensory prior information is encoded in the cross-talks in a distributed manner. The most striking discovery is that the cross-links and the reciprocal couplings form an antagonistic pair and play complementary roles. Our results reveal the crucial dependence of the optimal network structure on the statistics of multisensory stimuli, especially the prior information.

Keywords: multisensory integration; decentralized model

Introduction

The neural system can integrate multisensory cues in a near optimal way, as predicted by Bayesian inference (Ernst & Banks, 2002; Gu, Angelaki, & DeAngelis, 2008; Angelaki, Gu, & DeAngelis, 2009). However, exactly how the brain implements optimal multisensory processing remains largely unknown. To realize multisensory Bayesian inference, it is essential for the neural estimator to access the information of all sensory inputs and utilize the prior knowledge about the sensory cues (Ma, Beck, Latham, & Pouget, 2006; Vilares & Körding, 2011). Recordings of neuronal activities when they are utilizing multisensory information provide evidence for the neural representation of probability distributions and the neural implementation of Bayesian inference. For instance, in the heading direction perception tasks, multisensory neurons that are tuned to both visual and vestibular cues have been found extensively in several cortical areas (Gu et al., 2008; Chen, DeAngelis, & Angelaki, 2011). Among these areas, there are abundant *cross-talks*, referring to both *cross-links* (which directly feed low-level inputs from one modality to another) and

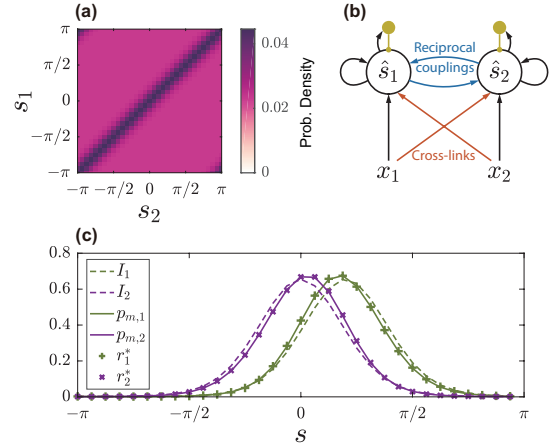


Figure 1: Bi-modality Bayesian inference. (a) The joint distribution of the congruent prior with $p_c = 0.1$ and $\kappa_p = 12$. (b) The bi-modular recurrent neural network model. The population activity in each module represents the neural estimate of the corresponding stimulus \hat{s}_i . Cross-talks between the modules consists of the cross-links (red arrows) and the reciprocal couplings (blue arrows). The yellow circles are the inhibitory pools. (c) The green (purple) dashed line is the example likelihood function of the stimulus s for modality 1(2) when $x_1 = 0$ ($x_2 = 0.63$). The green (purple) solid line is the marginal posterior of modality 1(2) according to Bayes' rule. The green plus signs (purple crosses) are the steady state firing rates in module 1(2) in the optimized network. $\kappa = 3$.

reciprocal links (which convey integrated signals among high-level areas) (Smith, Greenlee, DeAngelis, & Angelaki, 2017). These areas are likely to form a decentralized architecture where cross-talks between multiple pathways at different levels play a crucial role in integrating multisensory information. It has been shown that, for a fully correlated prior, the reciprocal links between multisensory modules encodes the correlation in the prior and the uncertainty of the posterior is represented in the sampling process of the neural dynamics (W.-H. Zhang, Chen, Rasch, & Wu, 2016). However, the functional role of the cross-links and the interplay between reciprocal links and cross-links remain unclear for the partially correlated prior in general.

In this work, we address the challenging issue of how the Bayes-optimal multisensory processing is implemented through the coordination of cross-links and reciprocal links. We formulate multisensory integration as a mathematical problem of optimizing network structure under the constraint that for a given stimulus prior, the network's output matches

the profile of the posterior of the stimulus. This is equivalent to requiring that the network realizes Bayesian inference when the sensory cues are sampled from their prior over many trials. Remarkably, distinct roles played by the cross-links and the reciprocal couplings are revealed in this framework. These results generate predictions about the structural pre-requisites for multisensory integration, which can be tested in future experiments and inspire novel artificial intelligent systems.

Multisensory Bayesian Inference with a Partially Correlated Prior

Take the bi-modality inference as an example. Let S_1 and S_2 denote the two external stimuli of different modalities. For convenience, we consider circular variables in the range $|s_i| \leq \pi$ for $i = 1, 2$.¹ The compound multisensory prior is composed of an independent part and a correlated part (Körding et al., 2007; Sato, Toyozumi, & Aihara, 2007),

$$p(s_1, s_2) = (1 - p_c) p(s_1) p(s_2) + p_c q(s_1, s_2). \quad (1)$$

Here, the first part of the prior is merely a product of the two marginal priors $p(s_1)$ and $p(s_2)$. The second part $q(s_1, s_2)$ describes the dependency between S_1 and S_2 when they originate from a common cause, with the probability of them coming from a common cause denoted by p_c . In this work, we assume that the marginal priors are flat [$p(s_1) = p(s_2) = (2\pi)^{-1}$], and the two stimuli are positively correlated in a way described by the von Mises distribution. This prior, which we refer to as the *congruent prior*, is given by $p(s_1, s_2) = [(1 - p_c) + p_c e^{\kappa_p \cos(s_1 - s_2)} / \mathcal{I}_0(\kappa_p)] / 4\pi^2$, where κ_p is the concentration of the correlated prior. An instance of this prior is shown in Fig. 1(a).

The corresponding sensory observations are X_1 and X_2 . We assume that unisensory likelihoods are independent von Mises distributions (Murray & Morgenstern, 2010), $p(x_1, x_2 | s_1, s_2) = \prod_i e^{\kappa_i \cos(s_i - x_i)} / 2\pi \mathcal{I}_0(\kappa_i)$ [dashed lines in Fig. 1(c)], where κ_i are the reliability of the sensory observation in modality i , and $\mathcal{I}_n(\cdot)$ is the modified Bessel function of the first kind and order n introduced to normalize the probability.

It is believed that Bayesian optimal perception in each modality requires marginalizing the posterior distribution over all other modalities (Körding et al., 2007; Beck, Latham, & Pouget, 2011). For convenience, we denote the marginal posterior $p(S_i = s | x_1, x_2)$ by $p_{m,i}(s | x_1, x_2)$ [solid lines in Fig. 1(c)]. The Bayesian estimate \hat{s}_i is the mean direction of the marginal posterior distribution, $\hat{s}_i = \arg \phi_i$, where $\phi_i \equiv \int_{-\pi}^{\pi} p_{m,i}(s | x_1, x_2) e^{i s} ds$, and i is the imaginary unit. The reliability of the Bayesian estimate is $\hat{\kappa}_i = \mathcal{A}^{-1}(|\phi_i|)$, where \mathcal{A}^{-1} is the inverse function of $\mathcal{A}(\kappa) \equiv \mathcal{I}_1(\kappa) / \mathcal{I}_0(\kappa)$ (Mardia & Jupp, 2000).

¹In this work, S_i, X_i are random variables, and s_i, x_i are instances of them.

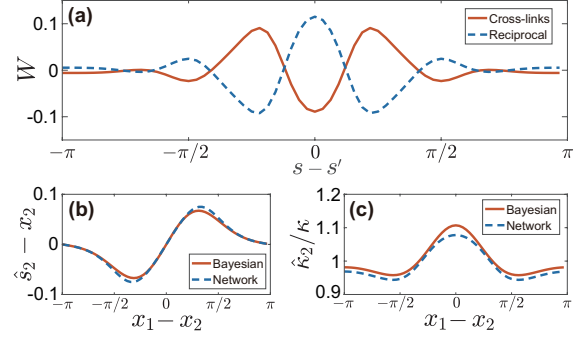


Figure 2: Optimized cross-talks. (a) The cross-links (the red solid line) and the reciprocal couplings (the blue dashed line) form an antagonistic pair. (b) The biases of the Bayesian (the red solid line) and the network (the blue dashed line) estimates. (c) The reliabilities of the Bayesian (the red solid line) and the network (the blue dashed line) estimates, scaled to the reliability of the likelihood function. Parameters: $\kappa = 3$, $\kappa_p = 12$, $p_c = 0.1$, and $k_1 \rho = 0.9$.

Two Coupled Modules Implements Optimal Multisensory Integration

We combine two modules of recurrent neural fields and introduce feedforward cross-links and reciprocal couplings to carry out the interaction between the two modules,

$$\tau \dot{U}_i(s, t) = -U_i(s, t) + \rho \int_{-\pi}^{\pi} ds' \sum_{j=1,2} \left[W_{ij}^{\text{ff}}(s, s') I_j(s'; x_j) + W_{ij}^{\text{rec}}(s, s') r_j(s', t) \right], \quad i = 1, 2, \quad (2)$$

where τ is the neuronal time scale, $U_i(s, t)$ is the synaptic input on the neuron preferring s in module i at time t , r_i is the firing rate of module i , I_i is the external input on module i , W_{ii}^{ff} is the same-side feedforward connection in module i , W_{ii}^{rec} is the same-side recurrent connection in module i , W_{ij}^{ff} ($i \neq j$) is the cross-links from module j to module i , and W_{ij}^{rec} ($i \neq j$) is the reciprocal couplings from module j to module i . This network architecture is illustrated in Fig. 1(b).

In the above equation, the firing rate r_i is related to the synaptic input through the activation function $r(s, t) = f(U(s, t))$. Recently, due to its success in accounting for important features of multisensory integration, such as the principle of inverse effectiveness and the spatial principle, divisive normalization has been proposed to be a canonical integration operation (Beck et al., 2011; Carandini & Heeger, 2012). Here, we incorporate the global inhibition, a mechanism to achieve divisive normalization, into the activation function $r_i(s, t) = [U_i(s, t)]_+^2 / (1 + k_1 \rho \int [U_i(s', t)]_+^2 ds')$, where k_1 is the strength of global inhibition, and $[x]_+$ is equal to x when $x \geq 0$, otherwise 0 (K. Zhang, 1996).

The inputs on the two modules are assumed to represent the likelihood functions, $I_i(s; x_i) = e^{\kappa_i \cos(s - x_i)} / 2\pi \mathcal{I}_0(\kappa_i)$. In the present study, we assume that both sensory observations

have the same reliability $\kappa_1 = \kappa_2$, and simply denote it as κ . Given these inputs, the network activities at the steady state $r_i^*(s; x_1, x_2)$ are assumed to represent the marginal posterior distributions (the dependence on x_1, x_2 will be omitted in the notation of r_i^* in the later part). We define the mean squared error between the stationary firing rate and the corresponding marginal posterior in both modules as the loss function $\bar{L} \equiv \langle \sum_{i=1,2} \int_{-\pi}^{\pi} ds \|r_i^*(s) - p_{m,i}(s|x_1, x_2)\|^2 \rangle_{x_1, x_2}$. By minimizing this loss function with respect to the coupling weights, one can find out what kind of couplings can facilitate the optimal computation of marginal posteriors when the external inputs on the two modules are the corresponding likelihood functions.

Network Optimization

As with most neural network models, the loss function can be minimized through stochastic gradient descent (SGD) (Wang, Zhang, Wong, & Wu, 2017). Although applicable to the general case, this method is slow and subject to several well known difficulties in training recurrent neural networks (Pascanu, Mikolov, & Bengio, 2013). Therefore, we seek theoretical results under certain assumptions, which can reveal important aspects of the networks that can be generalized.

We developed a set of basis functions derived from the von Mises function to facilitate perturbative theoretical analysis for non-flat and partially correlated priors. An analogous approach had been used in neural field models considering Gaussian-shaped profiles of neural activities in an extended domain (Fung, Wong, & Wu, 2008). Here, in the case of a circular domain, the set of basis functions describes different dynamical modes of the neural activity profile perturbed from the von Mises function. For the weakly correlated congruent prior (small p_c), we can decompose the coupling weights in terms of these basis functions and perform gradient descent with respect to their coefficients. Since these basis functions efficiently capture the features of the coupling weights, this method not only enjoys speed, but also provide insights into the correspondence between the neural structure and the statistics of stimuli.

Antagonistic cross-talks

An example of the optimized network and its performance are shown in Fig. 2. We find that the cross-links and the reciprocal couplings are antagonistic to each other.

The reciprocal coupling is excitatory in the short range but inhibitory at the flank [the blue dashed line in Fig. 2(a)]. The narrow excitatory center and the broad inhibitory surround resemble the Mexican-hat couplings well studied in the neural field theory (Bressloff, 2012). They serve to stabilize a more concentrated, thus more reliable population representation in both modules (W.-H. Zhang et al., 2016).

The cross-link behaves in an entirely opposite manner, which is inhibitory in the short range, while excitatory in the intermediate range, and gradually decays with minor ripples in the long range [the red solid line in Fig. 2(a)]. This inhibition in the short range is likely to cancel out redundant in-




Figure 3: Optimized cross-talks with weight decay. (a-d) The weight-decay factor for cross-links (the blue solid lines) λ_{cross} is 2, 8, 32, and 128, respectively. The weight-decay factor for reciprocal couplings (the red dashed lines) $\lambda_{\text{recip}} = 0$. (e-h) λ_{recip} is 2, 8, 32 and 128, respectively, and $\lambda_{\text{cross}} = 0$. Parameters are the same as those in Fig. 2.

formation when the two channels carry identical information. The broader excitatory region in the surround is advantageous for allowing the integration of two cues with moderate disparity. The SGD-optimized couplings with the opposite prior also show similar antagonistic patterns (data not shown).

The steady state firing rates of the optimized network are compared with the marginal posteriors in Fig. 1(c) for $x_1 - x_2 = 0.63$. The network's behavior at the steady state seamlessly captures the attraction between the two marginal posteriors due to the congruent prior. For general values of the disparity $x_1 - x_2$, the biases and reliabilities of the networks' estimates share the same trend as those of the Bayesian optimal estimates with minor discrepancies [Fig. 2(b) and (c)].

Structural robustness

To assess the robustness of the network and the different roles played by the two types of cross-talks, we apply weight-decay regularization on either of them. λ_{cross} and λ_{recip} denote the weight-decay factors on the L^2 -norm of the cross-links and the reciprocal couplings, respectively.

When $\lambda_{\text{recip}} = 0$, as λ_{cross} increases, both of the cross-links and the reciprocal couplings are weakened, while maintaining their general center-surround profile [Fig. 3(a)-(d)]. When $\lambda_{\text{cross}} = 0$, as λ_{recip} increases, the reciprocal coupling is suppressed progressively, while maintaining its profile [red dashed lines in Fig. 3(e)-(h)]. However, for large values of λ_{recip} , the cross-links are completely flipped [Fig. 3(g) and (h)]. This indicates the importance of the reciprocal couplings for a robust network structure, in the sense that a stringent constraint on them may need to be compensated by substantial changes in the cross-links.

We further investigate how the coupling profiles change with the characteristics of the prior. The antagonistic pair of cross-links and reciprocal couplings undergoes minor changes to accommodate a range of p_c at a fixed κ_p and κ [Fig. 4(a) and (b)], showing remarkable robustness. The con-




Figure 4: The optimized cross-talks for different likelihoods and priors. The color scale denotes the strength of the cross-links in (a) and (c), and the reciprocal couplings in (b) and (d). In (a) and (b), $p_c \in (0, 0.2]$, and $\kappa_p = 12$. In (c) and (d), $\kappa_p \in (0, 13]$, and $p_c = 0.1$. For (a)-(d), $\kappa = 3$.

centration of the correlation between the cues κ_p generally defines the range within which two cues should be integrated. Therefore, it naturally determines the widths of the cross-links and the reciprocal coupling [Fig. 4(c) and (d)].

Conclusion and Discussion

We have developed a theoretical framework to link the network structure of the multisensory brain region to the statistics of Bayesian inference. Information about the priors is stored in the cross-links and the reciprocal couplings, which form an antagonistic pair and play complementary roles in the optimal multisensory inference. Specifically, the reciprocal couplings contribute to the integration of correlated information and the robustness of the network structure, while the cross-links cancels out the uncorrelated component in the raw inputs and improve integration for cues with moderate disparity.

The optimal structure we found has implications to the decentralized architecture for multisensory information processing. We have demonstrated how the prior information affects the network circuitry between multisensory brain regions, which can be compared with future neurophysiological data as an evidence for the decentralized architecture. Furthermore, the optimal representation of the posterior distribution will be able to facilitate causal inference in a downstream layer, and this perturbative theoretical framework will lay the foundation for our future explorations on causal inference (Shams & Beierholm, 2010; Körding et al., 2007).

Acknowledgments

This work is supported by the Research Grants Council of Hong Kong (16322616 and 16306817), National Basic Research Program of China (2014CB846101), and the NSF China (31261160495).

References

Angelaki, D. E., Gu, Y., & DeAngelis, G. C. (2009). Multisensory integration: psychophysics, neurophysiology, and computation. *Curr Opin Neurobiol*, *19*(4), 452-458.

Beck, J. M., Latham, P. E., & Pouget, A. (2011). Marginalization in neural circuits with divisive normalization. *J Neurosci*, *31*(43), 15310-15319.

Bressloff, P. C. (2012). Spatiotemporal dynamics of continuum neural fields. *J Phys A: Math Theor*, *45*(3), 033001.

Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nat Rev Neurosci*, *13*(1), 51-62.

Chen, A., DeAngelis, G. C., & Angelaki, D. E. (2011). A comparison of vestibular spatiotemporal tuning in macaque parietoinsular vestibular cortex, ventral intraparietal area, and medial superior temporal area. *J Neurosci*, *31*(8), 3082-3094.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429-433.

Fung, C. C. A., Wong, K. Y. M., & Wu, S. (2008). Dynamics of neural networks with continuous attractors. *Europhys Lett*, *84*(1), 18002.

Gu, Y., Angelaki, D. E., & DeAngelis, G. C. (2008). Neural correlates of multisensory cue integration in macaque MSTd. *Nat Neurosci*, *11*(10), 1201-1210.

Körding, K. P., Beierholm, U., Ma, W. J., Quartz, S., Tenenbaum, J. B., & Shams, L. (2007). Causal inference in multisensory perception. *PLOS ONE*, *2*(9), e943.

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat Neurosci*, *9*(11), 1432-1438.

Mardia, K. V., & Jupp, P. E. (2000). *Directional statistics*. John Wiley & Sons Ltd, Chichester, UK.

Murray, R. F., & Morgenstern, Y. (2010). Cue combination on the circle and the sphere. *J Vision*, *10*(11), 15.

Pascanu, R., Mikolov, T., & Bengio, Y. (2013). On the difficulty of training recurrent neural networks. *JMLR W&CP*, *28*(3), 1310-1318.

Sato, Y., Toyozumi, T., & Aihara, K. (2007). Bayesian inference explains perception of unity and ventriloquism aftereffect: Identification of common sources of audiovisual stimuli. *Neural Comput*, *19*(12), 3335-3355.

Shams, L., & Beierholm, U. R. (2010). Causal inference in perception. *Trends Cogn Sci*, *14*(9), 425-432.

Smith, A. T., Greenlee, M. W., DeAngelis, G. C., & Angelaki, D. E. (2017). Distributed visual-vestibular processing in the cerebral cortex of man and macaque. *Multisens Res*, *30*(2), 91-120.

Vilares, I., & Körding, K. P. (2011). Bayesian models: The structure of the world, uncertainty, behavior, and the brain. *Ann NY Acad Sci*, *1224*(1), 22-39.

Wang, H., Zhang, W.-H., Wong, K. Y. M., & Wu, S. (2017). How the prior information shapes neural networks for optimal multisensory integration. In *14th International Symposium on Neural Networks, ISNN 2017, Sapporo, Japan*.

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *J Neurosci*, *16*(6), 2112-2126.

Zhang, W.-H., Chen, A., Rasch, M. J., & Wu, S. (2016). Decentralized multisensory information integration in neural systems. *J Neurosci*, *36*(2), 532-547.