



Canonical cortical circuits and the duality of Bayesian inference and optimal control

Kenji Doya

The duality of sensory inference and motor control has been known since the 1960s and has recently been recognized as the commonality in computations required for the posterior distributions in Bayesian inference and the value functions in optimal control. Meanwhile, an intriguing question about the brain is why the entire neocortex shares a canonical six-layer architecture while its posterior and anterior halves are engaged in sensory processing and motor control, respectively.

Here we consider the hypothesis that the sensory and motor cortical circuits implement the dual computations for Bayesian inference and optimal control, or perceptual and value-based decision making, respectively. We first review the classic duality of inference and control in linear quadratic systems and then review the correspondence between dynamic Bayesian inference and optimal control. Based on the architecture of the canonical cortical circuit, we explore how different cortical neurons may represent variables and implement computations.

Address

Neural Computation Unit, Okinawa Institute of Science and Technology Graduate University, 1919-1 Tancha, Onna, Okinawa 904-0495, Japan

Corresponding author: Doya, Kenji (doya@oist.jp)

Current Opinion in Behavioral Sciences 2021, **41**:160–166

This review comes from a themed issue on **Value-based decision-making**

Edited by **Laura Bradfield** and **Bernard Balleine**

For a complete overview see the [Issue](#) and the [Editorial](#)

Available online 17th September 2021

<https://doi.org/10.1016/j.cobeha.2021.07.003>

2352-1546/© 2021 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Introduction

Sensory perception and motor control are the most fundamental functions of the brain. Although they are often studied separately, sensory perception and motor control are dependent on each other, which calls for an integrated approach. In addition to composing a sensory-motor loop, computations for sensory inference and optimal control have been shown to have a similarity. Rudolph Kalman [1] showed that the equations used for optimal sensory inference by the Kalman filter are similar to the equations used for optimal motor control by a linear quadratic regulator (LQR) [2]. This is known as Kalman's duality.

More recently, researchers in reinforcement learning discovered more general correspondences between the computations for the posterior distribution in dynamic Bayesian inference and the value function in reinforcement learning [3–10,11**]. This notion has created a research strategy called control as inference, or reinforcement learning as inference, which provides novel mathematical insights and helps the development of new reinforcement learning algorithms.

Regarding the brain's architecture for sensory perception and motor control, the most fundamental division is between the posterior half of the cerebral cortex, which is mostly involved in sensory perception, and the anterior half, which is mostly involved in motor control, planning and decisions. An interesting unanswered, or unasked, question is: Why does the entire cerebral cortex share the same canonical circuit architecture [12,13] characterized by a common six-layer architecture with specific types of neurons and connections for implementing sensory inference and motor control?

Here we consider the hypothesis that the sensory and motor cortical circuits evolved to make the dual computations for sensory inference and optimal control, or perceptual and value-based decision making, respectively. We first review the classic duality of inference and control in linear quadratic systems, known as Kalman's duality, and then review a more general correspondence between dynamic Bayesian inference and optimal control. We then explore how different types of cortical neurons in sensory and motor cortices may represent different variables and what cortical dynamics may realize required computations. We further discuss what experimental and computational approaches are required for scrutinizing this dual cortical circuit hypothesis.

The duality of control and inference

Kalman filter [1] is the standard method for keeping track of the signal of interest despite noisy observation based on the assumption of linear dynamics and Gaussian noise (Box 1). Kalman pointed out that the set of equations for updating the estimates of the mean and the covariance of the state variable has the same structure as the equations for optimal control of a linear dynamical system with Gaussian noise [2]. This is known as Kalman's duality.

While Kalman's duality has been documented in textbooks of control theory and signal processing as a matter of mathematical beauty, recent researchers in reinforcement

Box 1 Kalman's duality (Table 1, Figure 1).

The computations for optimal filtering and optimal control under linear Gaussian assumptions reduces to solving the same form of matrix equations [1,7].

Figure 1: A. Kalman filter. We consider a linear discrete-time dynamical system

$$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \mathbf{v}_t \quad (1)$$

$$\mathbf{y}_t = C\mathbf{x}_t + \mathbf{w}_t \quad (2)$$

where \mathbf{x}_t is the state, \mathbf{u}_t is the action input, \mathbf{y}_t is the sensory observation, and \mathbf{v}_t and \mathbf{w}_t are the state and observation noises with covariance matrices S and U , respectively. We aim to estimate the changing state $\mathbf{x}_1, \dots, \mathbf{x}_t$ iteratively from the observations $\mathbf{y}_1, \dots, \mathbf{y}_t$. We represent the uncertainty of the state by a Gaussian distribution $\mathbf{x}_t \sim \mathcal{N}(\bar{\mathbf{x}}_t, \Sigma_t)$.

With the state transition (1), the mean and the variance of the state distribution evolve as

$$\hat{\mathbf{x}}_{t+1} = A\bar{\mathbf{x}}_t + B\mathbf{u}_t \quad (3)$$

$$\hat{\Sigma}_{t+1} = S + A\Sigma_t A' \quad (4)$$

where A' is the transpose matrix of A . With a new observation \mathbf{y}_{t+1} , the distribution is updated by Bayesian inference with the predicted distribution $\mathcal{N}(\hat{\mathbf{x}}_{t+1}, \hat{\Sigma}_{t+1})$ as the prior and the observation giving the likelihood $\mathcal{N}(\mathbf{y}_{t+1} - C\hat{\mathbf{x}}_{t+1}; 0, U)$. This leads to an update of the mean of the state distribution in proportion to the prediction error

$$\bar{\mathbf{x}}_{t+1} = \hat{\mathbf{x}}_{t+1} + K_t(\mathbf{y}_{t+1} - C\hat{\mathbf{x}}_{t+1}) \quad (5)$$

where the update gain is given by

$$K_t = A\Sigma_t C'(U + C\Sigma_t C')^{-1} \quad (6)$$

This is called the filter gain, or Kalman gain, which becomes large when the state uncertainty $\Sigma(t)$ is large. The variance of the state distribution is also updated by the Kalman gain as

$$\Sigma_{t+1} = \hat{\Sigma}_{t+1} - K_t C \Sigma_t A' \quad (7)$$

which generally reduces the uncertainty. Taken together (4), (6) and (7), the update of the state covariance is given as

$$\Sigma_{t+1} = S + A\Sigma_t A' - A\Sigma_t C'(U + C\Sigma_t C')^{-1} C\Sigma_t A' \quad (8)$$

B: Linear quadratic regulator (LQR). Here we consider the same dynamical system (1) and

consider the cost function (negative reward)

$$l(\mathbf{x}, \mathbf{u}) = \frac{1}{2} \mathbf{x}' Q \mathbf{x} + \frac{1}{2} \mathbf{u}' R \mathbf{u} \quad (9)$$

where Q and P are matrices defining the state and action costs [2]. The aim is to minimize the cumulative cost

$$J = \sum_{t=1}^T l(\mathbf{x}_t, \mathbf{u}_t) \quad (10)$$

In the linear Gaussian assumption, the state value function $V(\mathbf{x}, t)$ can be represented by a quadratic form of the state vector \mathbf{x} and a matrix P_t as

$$V(\mathbf{x}, t) = \frac{1}{2} \mathbf{x}' P_t \mathbf{x} = E \left[\sum_{s=t}^T l(\mathbf{x}_s, \mathbf{u}_s) \mid \mathbf{x}_t = \mathbf{x} \right] \quad (11)$$

The optimal action is given by the Bellman equation

$$\frac{1}{2} \mathbf{x}' P_t \mathbf{x} = \min_{\mathbf{u}} \left[l(\mathbf{x}, \mathbf{u}) + \frac{1}{2} (A\mathbf{x} + B\mathbf{u})' P_{t+1} (A\mathbf{x} + B\mathbf{u}) \right] \quad (12)$$

The solution leads to a feedback control law

$$\mathbf{u}(t) = -G_t \mathbf{x}_t \quad (13)$$

where the feedback control gain is given by

$$G_t = (R + B' V_{t+1} B)^{-1} B' P_{t+1} A \quad (14)$$

The matrix for the value function is computed backward in time

$$P_t = Q + A' P_{t+1} A - A' P_{t+1} B G_t \quad (15)$$

with the terminal condition of $P_T = Q$. From (14) and (15), the update equation for V is

$$P_t = Q + A' P_{t+1} A - A' P_{t+1} B (R + B' P_{t+1} B)^{-1} B' P_{t+1} A \quad (16)$$

which has a similar form with the equation (8) for Kalman filter.

The table below summarizes the correspondence between Kalman filter and LQR [1,7,8].

Figure 1

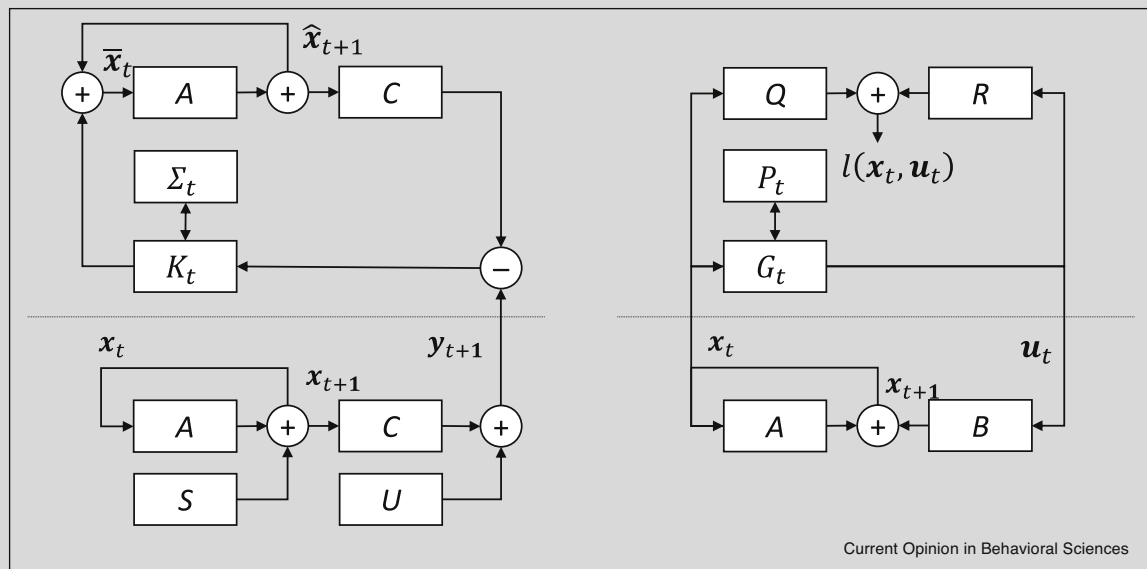


Table 1

The duality of the Kalman filter [1] and the linear quadratic regulator (LQR) [2]

Kalman filter	Linear quadratic regulator
Covariance of state distribution Σ_t	Coefficients of quadratic state value P_t
Filter gain $K_t = A\Sigma_t C' (U + C\Sigma_t C')^{-1}$	Control gain $G_t = (R + B'P_{t+1}B)^{-1}B'P_{t+1}A$
State dynamics A	Transpose state dynamics A'
Observation C	Action sensitivity B'
State noise covariance S	State cost coefficients Q
Observation noise covariance U	Action cost coefficients R
Forward from $\Sigma_1 = S$	Backward from $P_T = Q$

learning found that the relationship can extend beyond linear-Gaussian cases and have developed novel reinforcement learning and control algorithms based on the notion.

Emanuel Todorov pointed out the general duality between computations needed for posterior distributions in dynamic Bayesian inference and the value functions in optimal control and reinforcement learning (Box 2) [7–9]. A similar correspondence was also formulated by other researchers as well [3–6,10].

Based on this notion, Todorov realized that, by defining the cost of action as the divergence of the state transition probability from that by ‘passive dynamics,’ the exponentiated state value function can be computed linearly, which drastically reduces the required data and computation, and enables compositionality of value functions for different goal rewards [14,15]. A similar approach was also derived by Kappen [3,4].

Most recently, Sergey Levine reviewed all these works and formulated a probabilistic graphical model (PGM)

having the optimality variable, which takes 1 if a state-action pair is optimal (Figure 2b) [11•]. By assuming that the optimality variable follows a probability given by the exponential of a reward function, a standard message passing algorithm for PGM turns into the update equations for state and action value functions. For this conversion to hold, the objective function should include a regularization term for the entropy of action policy, which is known as maximum-entropy reinforcement learning [16].

Table 2 summarizes the correspondence between the components of Bayesian inference and optimal control. The framework presents a unified theoretical basis for efficient and robust reinforcement learning algorithms, such as the soft actor-critic [17•], and is expected to promote derivation of novel algorithms.

Canonical cortical circuits

The cerebral neocortex has a common six-layer architecture, known as the *canonical cortical circuit* (Figure 3) [12,13,18]. While most studies focused on the sensory

Box 2 Duality of Bayesian inference and optimal control [50,51].

Instead of just formal similarity, there is a meaningful correspondence between dynamic Bayesian inference and optimal control or reinforcement learning [7-9].

Figure 2: A. Dynamic Bayesian inference. Here we consider the stochastic state dynamics

$$p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) \quad (1)$$

and the sensory observation

$$p(\mathbf{o}_t|\mathbf{s}_t) \quad (2)$$

In active perception, the aim is to estimate the state trajectory behind a series of actions and observations

$p(\mathbf{s}_1, \dots, \mathbf{s}_T|\mathbf{a}_1, \mathbf{o}_1, \dots, \mathbf{a}_T, \mathbf{o}_T)$. In observation learning, the sequence of actions, such as muscle forces, are estimated from observation as $p(\mathbf{s}_1, \mathbf{a}_1, \dots, \mathbf{s}_T, \mathbf{a}_T|\mathbf{o}_1, \dots, \mathbf{o}_T)$.

Here we outline the simplest case involving no action, $p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) = p(\mathbf{s}_{t+1}|\mathbf{s}_t)$. This is an example of a hidden Markov model (HMM) and a standard algorithm for solving the problem is the forward-backward algorithm [50,51]. We define the forward message as the joint probability of the observations up to t and the state \mathbf{s}_t

$$\alpha(\mathbf{s}_t) = p(\mathbf{o}_1, \dots, \mathbf{o}_t, \mathbf{s}_t) \quad (3)$$

and the backward message as the conditional probability of observations after the state \mathbf{s}_t

$$\beta(\mathbf{s}_t) = p(\mathbf{o}_{t+1}, \dots, \mathbf{o}_T|\mathbf{s}_t) \quad (4)$$

The forward and backward messages are iteratively computed as

$$\alpha(\mathbf{s}_t) = p(\mathbf{o}_t|\mathbf{s}_t) \int_{\mathcal{S}} \alpha(\mathbf{s}_{t-1}) p(\mathbf{s}_t|\mathbf{s}_{t-1}) d\mathbf{s}_{t-1} \quad (5)$$

$$\beta(\mathbf{s}_t) = \int_{\mathcal{S}} \beta(\mathbf{s}_{t+1}) p(\mathbf{o}_{t+1}|\mathbf{s}_{t+1}) p(\mathbf{s}_{t+1}|\mathbf{s}_t) d\mathbf{s}_{t+1} \quad (6)$$

The posterior probability is then computed from the product of the two messages.

$$p(\mathbf{s}_t|\mathbf{o}_1, \dots, \mathbf{o}_T) \propto \alpha(\mathbf{s}_t) \beta(\mathbf{s}_t) \quad (7)$$

B. Optimal control as inference. There are several ways to cast optimal control problems into dynamic Bayesian inference problems, which facilitate common understanding and development of new algorithms [3-10].

One way is to consider an “optimality variable \mathcal{O}_t ” [11] which takes 1 when the state-action pair is optimal and assume that the reward function represents the log probability for the state-action pair to be optimal

$$p(\mathcal{O}_t = 1|\mathbf{s}_t, \mathbf{a}_t) = \exp(r(\mathbf{s}_t, \mathbf{a}_t)) \quad (8)$$

In this formulation, the posterior probability of state-action trajectory conditioned on the optimality variable to be 1 is given as

$$\begin{aligned} & p(\mathbf{s}_1, \mathbf{a}_1, \dots, \mathbf{s}_T, \mathbf{a}_T|\mathcal{O}_1 = 1, \dots, \mathcal{O}_T = 1) \\ & \propto p(\mathbf{s}_1) \prod_{t=1}^T p(\mathcal{O}_t = 1|\mathbf{s}_t, \mathbf{a}_t) p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) \\ & = \left\{ p(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) \right\} \left\{ \exp \sum_{t=1}^T r(\mathbf{s}_t, \mathbf{a}_t) \right\} \end{aligned} \quad (9)$$

which means that state-action trajectories that is feasible under the state dynamics with higher cumulative rewards have high posterior probability.

This posterior probability can be computed by backward message passing. We define backward messages representing the state-action pair or the state is optimal for time t and future

$$\beta_t(\mathbf{s}_t, \mathbf{a}_t) = p(\mathcal{O}_1 = 1, \dots, \mathcal{O}_T = 1|\mathbf{s}_t, \mathbf{a}_t) \quad (10)$$

$$\begin{aligned} \beta(\mathbf{s}_t) &= p(\mathcal{O}_1 = 1, \dots, \mathcal{O}_T = 1|\mathbf{s}_t) \\ &= \int_{\mathcal{A}} \beta_t(\mathbf{s}_t, \mathbf{a}_t) p(\mathbf{a}_t|\mathbf{s}_t) d\mathbf{a}_t \end{aligned} \quad (11)$$

where $p(\mathbf{a}_t|\mathbf{s}_t)$ is a policy prior, such as uniform action selection. These messages are computed backward in time as

$$\beta_t(\mathbf{s}_t, \mathbf{a}_t) = \int_{\mathcal{S}} \beta_{t+1}(\mathbf{s}_{t+1}) p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) \exp(r(\mathbf{s}_t, \mathbf{a}_t)) d\mathbf{s}_{t+1} \quad (12)$$

Then the optimal policy is derived from these messages as

$$p(\mathbf{a}_t|\mathbf{s}_t, \mathcal{O}_t = 1, \dots, \mathcal{O}_T = 1) \propto \frac{\beta_t(\mathbf{s}_t, \mathbf{a}_t)}{\beta_t(\mathbf{s}_t)} \quad (13)$$

These messages have correspondence with the value functions as

$$Q(\mathbf{s}_t, \mathbf{a}_t) = \log \beta_t(\mathbf{s}_t, \mathbf{a}_t) \quad (14)$$

$$V(\mathbf{s}_t) = \log \beta_t(\mathbf{s}_t) \quad (15)$$

under the addition of entropy-based regularization of the policy.

Figure 2




Table 2

The correspondence of dynamic Bayesian inference and optimal control

Bayesian inference	Optimal control
Posterior state distribution $p(\mathbf{s}_t \mathbf{o}_1, \dots, \mathbf{o}_T)$	State value function $\exp(V(\mathbf{s}_t))$
State dynamics $p(\mathbf{s}_{t+1} \mathbf{s}_t, \mathbf{a}_t)$	State dynamics $p(\mathbf{s}_{t+1} \mathbf{s}_t, \mathbf{a}_t)$
Observation model $p(\mathbf{o}_t \mathbf{s}_t)$	Reward function $p(\mathbf{o}_t = 1 \mathbf{s}_t, \mathbf{a}_t) = \exp(r(\mathbf{s}_t, \mathbf{a}_t))$
Forward message $\alpha(\mathbf{s}_t)$	–
Backward message $\beta(\mathbf{s}_t)$	Action value function $\exp(Q(\mathbf{s}_t, \mathbf{a}_t))$

cortex, the architecture of motor cortex with the thalamic inputs originating from the cerebellum and the basal ganglia has also been worked out [19]. A marked difference between the sensory and motor cortices is the thickness of the layer 4, which is densely populated by excitatory stellate cells in the sensory cortex. Despite a quantitative difference across areas, the basic architecture is preserved: layer 4 receives bottom-up thalamic and cortical input and projects to layers 2/3, where neurons have dense recurrent connections. Layer 2/3 pyramidal neurons project to higher cortical areas and also send output to layers 5/6. Layer 5 pyramidal neurons project to the cerebellum and the basal ganglia and layer 6 pyramidal neurons project to the thalamus and lower cortical areas.

Given the evidence suggesting that sensory perception is based on Bayesian inference combining top-down prediction and bottom-up sensory evidence [20–22], hypotheses have been proposed about how Bayesian inference can be implemented in the canonical cortical circuits [23–30]. Models of the motor cortex have also been proposed based on optimal control theory [31,32].

Dual cortical computation hypothesis

Considered together, the duality of Bayesian inference and optimal control and the canonical cortical circuits in the sensory and motor areas suggest that common computations for inference and control are implemented in


the common architecture of the neural circuits in the sensory and motor cortices, or the posterior and anterior halves of the cerebral cortex.

Figure 3 and Table 3 illustrate the canonical cortical circuits in the sensory and motor cortical areas and hypothetical representations of key variables of Bayesian inference and optimal control.

For sensory inference, the thalamic and cortical bottom-up inputs represent sensory observation \mathbf{o}_t , which is used for evaluating the likelihood $p(\mathbf{o}_t | \mathbf{s})$ of different world states \mathbf{s} in layer 4. The layer 2/3 combines the sensory likelihood with the predicted state $p(\mathbf{s}_t | \mathbf{s}_{t-1})$ to represent the surprise signal $p(\mathbf{o}_t | \mathbf{s}_t)$, which is sent to higher cortical areas. This information is also sent to layer 5/6 to update the posterior probability $p(\mathbf{s}_t | \mathbf{o}_1, \dots, \mathbf{o}_t)$. These computations may also be conditional on the top-down contextual signal \mathbf{z}_t , including the executed action \mathbf{a}_{t-1} , from higher cortical areas.

In optimal control, the reward function $r(\mathbf{s}, \mathbf{a})$ and the state value function $V(\mathbf{s})$ correspond to the log likelihood and log posterior probability in sensory inference, so that they would be represented by layer 4 and layer 5/6 neurons, respectively. The update of action value function $Q(\mathbf{s}, \mathbf{a})$ requires state transition model $p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$ and reward information, which is likely to be represented by layer 2/3 neurons. The action policy $p(\mathbf{a} | \mathbf{s})$ is

Figure 3



Canonical cortical circuits in sensory and motor cortices and a hypothetical realization of dynamic Bayesian inference and optimal control. **(a)** Possible realization of dynamic Bayesian inference in the sensory cortex. **(b)** Possible realization of optimal control in the motor cortex.

Table 3

Correspondences of dynamic Bayesian inference and optimal control, and their possible implementation in the canonical cortical circuit

Inference	Cortex	Control
Top-down signal z_t	L1 input	Top-down activation signal
Bottom-up signal $p(o_t s_t)$	L2/3 output	Action value $Q(s, a)$
Predictive model $p(s_t s_{t-1})$	L2/3 connection	Predictive model $p(s_{t+1} s_t, a_t)$
Bottom-up signal o_t	L4 input	Optimality signal o_t
Likelihood $p(o_t s)$	L4 output	Reward function $r(s, a)$
Posterior $p(s_t o_1, \dots, o_t)$	L5 output	State value $V(s)$
Top-down signal s_t	L6 output	Action $p(a_t s_t)$

computed by subtracting the state value from the action value, so that action may be selected in layer 5 or 6 and sent to lower cortical and subcortical areas. Note that the above is just one hypothetical realization and many other mappings of different roles to neurons and connections are conceivable.

There are many interesting open questions about the cortical implementation of the dual computations for Bayesian inference and optimal control. First, how the backward computation is realized in real time? In the visual cortex, evidence suggests that the alpha rhythm around 10 Hz carries top-down feedback information [33]

and underlies multi-modal sensory arbitration [34]. In the motor cortex, the beta rhythm around 20 Hz shows responses before execution or during imagination of movements [35,36]. These might be the correlates of periodic execution of backward computation.

Another important question is how the state transition model $p(s_{t+1}|s_t, a_t)$ and the sensory observation model $p(o_t|s_t)$ are learned, together with the internal representations of state s and action a . The roles of the cerebellar and the basal ganglia inputs through the thalamus to the motor cortex in learning is also an interesting question [37,38].

Finally, how are the parameters for Bayesian inference and optimal control regulated, such as the time frame of planning and the prior uncertainty of the state dynamics and sensory observation? The roles of neuromodulators, such as serotonin, noradrenaline and acetylcholine have been suggested [39–43,44[•]].

Given recent advances in two-photon calcium imaging [45,46[•]] and electrode array recording [47], it is now feasible to test such hypotheses regarding the implementation of Bayesian inference and optimal control by large-scale measurement of neural activities during sensorimotor tasks [38,48,49]. The correspondence between variables and mappings for Bayesian inference and optimal control as depicted in Table 3 may provide a basis for interpreting data from both sensory and motor cortices and coming up with a unified theory of cortical computation.

Conclusion

This article reviewed the duality between sensory inference and motor control and presented a hypothesis that the canonical posterior and anterior cortical circuits perform such dual computation. The author believes that this overreaching hypothesis is worthy of experimental testing by utilizing multi-area, multi-layer neural recording technologies. In addition to the basic operations for inference and control, the regulatory mechanisms for such computations and possible malfunctioning of such mechanisms would provide better understanding of the roles of neuromodulators and the causes of psychiatric disorders.

Conflict of interest statement

Nothing declared.

Acknowledgements

This work was supported by JSPS KAKENHI Grants 23120007, 16K21738, 16H06561 and 16H06563, AMED Grants JP19dm0207001h0006 and JP21dm0307009, and research support of Okinawa Institute of Science and Technology Graduate University to KD. The author thanks Takuya Isomura and the reviewer for their comments on the manuscript.

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
- of outstanding interest

1. Kalman RE: **A new approach to linear filtering and prediction problems**. *Trans ASME* 1960, **82-D**:35–45.
2. Kalman RE, Koepcke RW: **Optimal synthesis of linear sampling control systems using general performance indexes**. *Trans ASME* 1958, **80**:1820–1826.
3. Kappen H: **Linear theory for control of nonlinear stochastic systems**. *Phys Rev Lett* 2005, **95**:200201.
4. Kappen HJ, Gómez V, Opper M: **Optimal control as a graphical model inference problem**. *Mach Learn* 2012, **87**:159–182.
5. Toussaint M, Storkey A: **Probabilistic inference for solving discrete and continuous state Markov Decision Processes**. *Proceedings of the 23rd International Conference on Machine Learning (ICML 2006)* 2006:945–952.
6. Toussaint M: **Robot trajectory optimization using approximate inference**. *Proceedings of the 26th Annual International Conference On Machine Learning (ICML2009)* 2009:1049–1056.
7. Todorov E: **Optimal control theory**. In *Bayesian Brain: Probabilistic Approach to Neural Coding and Learning*. Edited by Doya K, Ishii S, Pouget A, Rao RP. MIT Press; 2007.
8. Todorov E: **General duality between optimal control and estimation**. *The 47th IEEE Conference on Decision and Control* 2008:4286–4292.
9. Todorov E: **Parallels between sensory and motor information processing**. In *The Cognitive Neurosciences*, edn 4. Edited by Gazzaniga MS. 2009.
10. Botvinick M, Toussaint M: **Planning as inference**. *Trends Cogn Sci* 2012, **16**:485–488.
11. Levine S: **Reinforcement learning and control as probabilistic inference: tutorial and review**. *arXiv* 2018 <https://arxiv.org/abs/1805.00909>
- A tutorial article on the correspondence of probabilistic inference and optimal control and reinforcement learning.
12. Felleman DJ, Van Essen DC: **Distributed hierarchical processing in the primate cerebral cortex**. *Cereb Cortex* 1991, **1**:1–47.
13. Douglas RJ, Martin KAC: **Canonical cortical circuits**. In *Handbook of Brain Microcircuits*, edn 2. Edited by Shepherd GM, Grillner S. Oxford University Press; 2017:15–21.
14. Todorov E: **Efficient computation of optimal actions**. *Proc Natl Acad Sci U S A* 2009, **106**:11478–11483.
15. Doya K: **How can we learn efficiently to act optimally and flexibly?** *Proc Natl Acad Sci U S A* 2009, **106**:11429–11430.
16. Ziebart B, Bagnell J, Dey A: **Modeling interaction via the principle of maximum causal entropy**. *International Conference on Machine Learning (ICML2010)* 2010:1247–1254.
17. Haarnoja T, Zhou A, Abbeel P, Levine S: **Soft actor-critic: off-policy maximum entropy deep reinforcement learning with a stochastic actor**. *arXiv* 2018 <https://arxiv.org/abs/1801.01290>
- One of the most successful reinforcement learning algorithms derived with the concept of control as inference.
18. Larkum M: **A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex**. *Trends Neurosci* 2013, **36**:141–151.
19. Kuramoto E, Furuta T, Nakamura KC, Unzai T, Hioki H, Kaneko T: **Two types of thalamocortical projections from the motor thalamic nuclei of the rat: a single neuron-tracing study using viral vectors**. *Cereb Cortex* 2009, **19**:2065–2077.
20. Knill DC, Pouget A: **The Bayesian brain: the role of uncertainty in neural coding and computation**. *Trends Neurosci* 2004, **27**:712–719.
21. Körding KP, Wolpert DM: **Bayesian integration in sensorimotor learning**. *Nature* 2004, **427**:244–247.
22. Doya K, Ishii S, Pouget A, Rao R: *Bayesian Brain: Probabilistic Approach to Neural Coding and Learning*. MIT Press; 2007.
23. Rao RPN, Ballard DH: **Dynamic model of visual recognition predicts neural response properties in the visual cortex**. *Neural Comput* 1997, **9**:721–763.
24. Lee TS, Mumford D: **Hierarchical Bayesian inference in the visual cortex**. *J Opt Soc Am A Opt Image Sci Vis* 2003, **20**:1434–1448.
25. Rao RP: **Bayesian computation in recurrent neural circuits**. *Neural Comput* 2004, **16**:1–38.
26. Friston K: **A theory of cortical responses**. *Philos Trans R Soc Lond B Biol Sci* 2005, **360**:815–836.
27. Ichisugi Y: **The cerebral cortex model that self-organizes conditional probability tables and executes belief**

- propagation.** *International Joint Conference on Neural Networks* 2007.
28. Deneve S: **Bayesian spiking neurons I: inference.** *Neural Comput* 2008, **20**:91-117.
 29. George D, Hawkins J: **Towards a mathematical theory of cortical micro-circuits.** *PLoS Comput Biol* 2009, **5**:e1000532.
 30. Bastos AM, Usrey WM, Adams RA, Mangun GR, Fries P, Friston KJ: **Canonical microcircuits for predictive coding.** *Neuron* 2012, **76**:695-711.
 31. Todorov E, Jordan MI: **Optimal feedback control as a theory of motor coordination.** *Nat Neurosci* 2002, **5**:1226-1235.
 32. Todorov E: **Optimality principles in sensorimotor control.** *Nat Neurosci* 2004, **7**:907-915.
 33. van Kerkoerle T, Self MW, Dagnino B, Gariel-Mathis MA, Poort J, van der Togt C, Roelfsema PR: **Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex.** *Proc Natl Acad Sci U S A* 2014, **111**:14332-14341.
 34. Minami S, Amano K: **Illusory jitter perceived at the frequency of alpha oscillations.** *Curr Biol* 2017, **27**:2344-2351.
 35. Wiesman AI, Koshy SM, Heinrichs-Graham E, Wilson TW: **Beta and gamma oscillations index cognitive interference effects across a distributed motor network.** *Neuroimage* 2020, **213**:116747.
 36. de Lange FP, Jensen O, Bauer M, Toni I: **Interactions between posterior gamma and frontal alpha/beta oscillations during imagined actions.** *Front Hum Neurosci* 2008, **2**:7.
 37. Doya K: **What are the computations of the cerebellum, the basal ganglia, and the cerebral cortex.** *Neural Netw* 1999, **12**:961-974.
 38. Tanaka YH, Tanaka YR, Kondo M, Terada SI, Kawaguchi Y, Matsuzaki M: **Thalamocortical axonal activity in motor cortex exhibits layer-specific dynamics during motor learning.** *Neuron* 2018, **100**:244-258.
 39. Doya K: **Metalearning and neuromodulation.** *Neural Netw* 2002, **15**:495-506.
 40. Yu AJ, Dayan P: **Uncertainty, neuromodulation, and attention.** *Neuron* 2005, **46**:681-692.
 41. Doya K: **Modulators of decision making.** *Nat Neurosci* 2008, **11**:410-416.
 42. Parr T, Friston KJ: **Uncertainty, epistemics and active inference.** *J R Soc Interface* 2017, **14**.
 43. Sales AC, Friston KJ, Jones MW, Pickering AE, Moran RJ: **Locus coeruleus tracking of prediction errors optimises cognitive flexibility: an active inference model.** *PLoS Comput Biol* 2019, **15**:e1006267.
 44. Doya K, Miyazaki KW, Miyazaki K: **Serotonergic modulation of cognitive computations.** *Curr Opin Behav Sci* 2021, **38**:116-123
This reviews recent advances in experimental approaches on the function of serotonin and proposes a novel hypothesis that serotonin signals the availability of time and resources, which affects parameters and strategies of decision making.
 45. Sofroniew NJ, Flickinger D, King J, Svoboda K: **A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging.** *eLife* 2016, **5**:e14472.
 46. Ota K, Oisi Y, Suzuki T, Ikeda M, Ito Y, Ito T, Uwamori H, Kobayashi K, Kobayashi M, Odagawa M *et al.*: **Fast, cell-resolution, contiguous-wide two-photon imaging to reveal functional network architectures across multi-modal cortical areas.** *Neuron* 2021, **109**:1810-1824
A recent two-photon imaging system allowing recording from 3mm square cortical areas.
 47. Steinmetz NA, Aydin C, Lebedeva A, Okun M, Pachitariu M, Bauza M, Beau M, Bhagat J, Böhm C, Broux M *et al.*: **Neuropixels 2.0: a miniaturized high-density probe for stable, long-term brain recordings.** *Science* 2021, **372**:eabf4588.
 48. Manita S, Suzuki T, Homma C, Matsumoto T, Odagawa M, Yamada K, Ota K, Matsubara C, Inutsuka A, Sato M *et al.*: **A top-down cortical circuit for accurate sensory perception.** *Neuron* 2015, **86**:1304-1316.
 49. Funamizu A, Kuhn B, Doya K: **Neural substrate of dynamic Bayesian inference in the cerebral cortex.** *Nat Neurosci* 2016, **19**:1682-1689.
 50. Baum LE: **An inequality and associated maximization technique in statistical estimation of probabilistic functions of Markov processes.** *Inequalities* 1972, **3**:1-8.
 51. Bishop CM: **Hidden Markov models.** In *Pattern Recognition and Machine Learning*. Edited by Bishop CM. Springer; 2006:610-635.