Title

Combination of successor representation-based appetitive learning and individual representation-based aversive learning is a good strategy and can be implemented in the cortico-basal ganglia circuits

Abbreviated title

Opponent learning with different representations

Kenji Morita[1,2*], Kanji Shimomura[1,3], & Yasuo Kawaguchi[4,5]

[1] Physical and Health Education, Graduate School of Education, The University of Tokyo, Tokyo 113-0033, Japan

[2] International Research Center for Neurointelligence (WPI-IRCN), The University of Tokyo, Tokyo 113-0033, Japan

[3] Department of Behavioral Medicine, National Institute of Mental Health, National Center of Neurology and Psychiatry, Kodaira 187-8551, Japan

[4] Brain Science Institute, Tamagawa University, Machida 194-8610, Japan

[5] National Institute for Physiological Sciences (NIPS), Okazaki 444-8787, Japan

**Corresponding author*

Kenji Morita, Ph.D. (ORCID iD: orcid.org/0000-0003-2192-4248)

Physical and Health Education, Graduate School of Education, The University of Tokyo

7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan

morita@p.u-tokyo.ac.jp

Author contributions

KM conceived of the possibility of differential representations in different cortico-basal ganglia pathways, and elaborated the idea in discussion with YK. KS considered to examine differential learning rates from positive and negative feedbacks, in discussion with KM. Based on these, KM developed the model presented in this manuscript, and conducted the simulations and analyses. KM and KS validated the program codes. KM drafted the manuscript, and revised it with inputs from KS and YK. According to the CRediT Taxonomy, Conceptualization: KM, KS, YK; Investigation: KM; Validation: KM, KS; Funding Acquisition: KM, YK; Writing – Original Draft Preparation: KM; Writing – Review & Editing: KM, KS, YK.

Conflict of interest

None

Acknowledgements

KM was supported by Grant-in-Aid for Scientific Research (No. 20H05049) of the Ministry of Education, Culture, Sports, Science and Technology in Japan (http://www.mext.go.jp/en/). YK was supported by Grant-in-Aid for Scientific Research (No. 17H06311, 20H03359) and by the Cooperative Study Program of National Institute for Physiological Sciences (Section of Electron Microscopy).

Manuscript information

23 pages, 8 figures, no table, no multimedia, no 3D models

Number of words: Abstract 250, Introduction 647, Discussion 1491

Abstract

While positive reward prediction errors (RPEs) and negative RPEs have equal impacts in the standard reinforcement learning, the brain appears to have distinct neural pathways for learning mainly from either positive or negative feedbacks, such as the direct and indirect pathways of the basal ganglia (BG). Given that distinct pathways may unevenly receive inputs from different neural populations and/or regions, how states or actions are represented can differ between the pathways. We explored whether combined use of different representations, coupled with different learning rates from positive and negative RPEs, has computational benefits. We considered an agent equipped with two learning systems, each of which adopted individual representation (IR) or successor representation (SR) of states. With varying the combination of IR or SR and also the learning rates from positive and negative RPEs in each system, we examined how the agent performed in a certain dynamic reward environment. We found that combination of SR-based system learning mainly from positive RPEs and IR-based system learning mainly from negative RPEs outperformed the other combinations, including IR-only or SR-only cases and the cases where the two systems had the same ratios of positive- and negative-RPE-based learning rates. In the best combination case, both systems show activities of comparable magnitudes with opposite signs, consistent with suggested profiles of BG pathways. These results suggest that particularly combining different representations with appetitive and aversive learning could be an effective learning strategy in a certain dynamic reward environment, and it might actually be implemented in the cortico-BG circuits.

Introduction

In the standard reinforcement learning (RL), updates based on positive reward prediction errors (RPEs) and those based on negative RPEs are executed in the same manner. However, in the brain, there appear to exist distinct neural circuits that mainly conduct either learning from positive feedbacks or learning from negative feedbacks. Specifically, a number of studies have suggested that the direct and indirect pathways of the basal ganglia (BG), originating from the striatal projection neurons (SPNs) expressing D1- and D2-type dopamine (DA) receptors (D1Rs and D2Rs), implement learning based primarily on positive and negative feedbacks, respectively (Frank et al., 2004; Hikida et al., 2010; Kravitz et al., 2012). There are also studies suggesting that distinct circuits involving the orbitofrontal cortex (OFC) operate for appetitive and aversive feedback-based learning (Groman et al., 2019). Computational works suggest that dual learning systems can realize estimation of costs and benefits (Collins and Frank, 2014; Möller and Bogacz, 2019), as well as estimation of not only the mean but also the uncertainty of rewards (Mikhael and Bogacz, 2016).

These existing dual learning-system models assume that both systems use the same way of representation of states or actions. Theoretically, various ways of representation can be considered, and different brain regions or neural populations may generally use different representations (Chen et al., 2014; Town et al., 2017; Wang et al., 2020). While there are evidences that the two BG pathways receive inputs from same types of corticostriatal neurons (Ballion et al., 2008; Kress et al., 2013), it has also been suggested that different neuron types (Lei et al., 2004; Reiner et al., 2010; Morita, 2014) and/or cortical areas (Wall et al., 2013) may not evenly target these pathways. As for the suggested distinct OFC circuits for appetitive and aversive learning, the suggested circuits are the amygdala-OFC and OFC-nucleus accumbens (NAc) pathways, respectively (Groman et al., 2019). Therefore, in both cases, it is conceivable that, between the appetitive and aversive learning systems, states/actions are represented by at least partially different neural regions or populations, and thus in different styles.

There are two largely different ways of state/action representation (Sutton and Barto, 1998). One is to represent each individual state/action separately. This simplest representation (or equivalent ones) has been explicitly or implicitly assumed to exist in the brain RL circuits in many previous studies. The other is to represent each state/action by a set of features (e.g., represent a point in a two-dimensional space by a set of x- and y-coordinates). Among various ways of feature-based representations, recent studies (Garvert et al., 2017; Momennejad et al., 2017; Russek et al., 2017; Stachenfeld et al., 2017; Russek et al., 2021) suggest that representation of states by their successors, named the successor representation (SR) (Dayan, 1993), may be adopted in the brain. SR contains information about state transitions in the environment under a given policy, and thereby enables the agent to quickly adapt to changes in distant outcomes through temporal difference (TD) RPE-based learning, without using an explicit model of the environment. It can thus beautifully explain (Russek

et al., 2017) the empirical suggestions that both (apparently) model-based behavior and model-free or habitual behavior are controlled by the DA-cortico-BG systems although different portions appear to be responsible (Balleine and O'Doherty, 2010; Dolan and Dayan, 2013) by assuming that SR and individual (punctate) representation (IR) are respectively used.

Given these suggestions and considerations, it seems possible that there exist two neural systems, which may differently learn from positive and negative RPEs, and may adopt different ways of state representations. In the present study, we examined its possible consequences through simulations. Specifically, we considered an agent equipped with two learning systems, and simulated its behavior in a reward navigation task in a certain dynamic reward environment, with the adopted representation, SR or IR, and the ratio of learning rates for positive versus negative TD-RPEs in each system freely and independently varied.

Methods

Simulated task

We simulated a reward navigation task in a certain dynamic reward environment. An agent was moving around in a 5 × 5 grid space (Figure 1A). The agent started from the fixed start state, which was location (1, 1), and moved to one of the neighboring states (two states at the corners, three states at the edges, and four states elsewhere) at each time step. During the initial 500 time steps, there was no reward, and the agent just moved around (Figure 1B). After that, a reward (always size 1) was introduced into one of the state (location) in the space. There were nine reward candidate states, where reward could potentially be placed, namely, (1, 5), (2, 5), (3, 5), (4, 5), (5, 1), (5, 2), (5, 3), (5, 4), and (5, 5) (i.e., the states on the two edges apart from the start state). During the next 500 time steps following the initial no-reward epoch, one of these nine candidate states was specified as the *special reward candidate state*, whereas the remaining eight candidate states were regarded as *normal candidate states*. There were in total nine rewarded epochs, each of which lasted for 500 time steps, and each one of the nine candidate states became the special candidate state in one of the nine epochs; the order was determined by pseudorandom permutation in each single simulation.

In the rewarded epochs, if the agent reached the rewarded state and obtained the reward, the agent was immediately carried back to the start state, and a new reward was introduced into a state, which was the special reward candidate state with 60% probability and one of the eight normal reward candidate states with an equal probability (i.e., 5%) each (Figure 1C).

We also examined variations of the task. In all the variations, there was the initial 500 time-steps no-reward epoch. In the first variation, the probability that reward was placed at the special candidate state, which was 60% in the original task, was varied to 70, 80, 90, or 100%. In the second variation, periodic resets of reward placement were introduced. Specifically, reward location (state) was reset at every 500, 250, 100, or 50 time steps in the rewarded epochs; at the reset timings, reward was placed at the special candidate state for the epoch with 60% and placed at one of the other (normal) candidate states with 5% each. In the third variation, the original nine rewarded epochs with different placements of special reward candidate state (500 time steps each) were abolished and instead there was only a single rewarded epoch with 4500 time steps. For this task variation, the special candidate state was varied only across simulations. In the fourth variation, reward was always placed at the special candidate state in each epoch and the duration of rewarded epoch was shortened from 500 time steps to 100 time steps while the number of rewarded epochs was increased from 9 to 45. The order of special reward candidate states was determined by five consecutive pseudorandom permutations of the nine candidate states. In the fifth variation, the probability of reward placement at the special candidate state was changed to 1/9 (11.11..%) so that reward was placed at each of the nine candidate states with equal probability (1/9). In the sixth variation, rewarded state was determined by a fixed order, namely,

(5, 1), (5, 5), (1, 5), and again (5, 1), and these were repeated throughout the task.

Model agent equipped with two learning systems

The agent's behavior was controlled by an RL model consisting of two learning systems (Figure 2). The two systems may use different ways of state representation. We considered the successor representation (SR) and the individual representation (IR) (also called the "punctate" representation), and examined the cases where system 1 and 2 used the SR and IR, respectively, and also the cases where only the SR or IR was used in both systems. Each system had its own *system-specific value* of each state (detailed below). The mean (average) of the system-specific values of the two systems was calculated and used as the *integrated value* of each state S, denoted as $V(S)$. At any states other than the rewarded state, the agent selected an action to move to one of the neighboring states depending on their integrated state values in a soft-max manner. Specifically, the agent selected the action to move to state S_i, among the neighboring states S_j, with probability

$$\exp(\beta V(S_i)) / \Sigma_j\{\exp(\beta V(S_j))\},$$

where β indicates the inverse temperature parameter, representing the degree of exploitation over exploration. When the agent was at state $S(t)$ at time-step t and moved to $S(t+1)$ at $t+1$, where $S(t)$ was not the rewarded state, TD-RPE was calculated based on the integrated state values:

$$\delta(t+1) = 0 + \gamma V(S(t+1)) - V(S(t)),$$

where γ indicates the time discount factor. When the agent was at $S(t)$, which was the rewarded state, TD-RPE

$$\delta(t+1) = 1 + 0 - V(S(t))$$

was calculated, again using the integrated state value.

In the system using the IR, the system-specific state value, $V_{system\text{-}i}(S(t))$, where i was 1 or 2, was directly updated based on the TD-RPE $\delta(t+1)$, with potentially different learning rates depending on whether the TD-RPE was positive (non-negative) or negative:

$$V_{system\text{-}i}(S(t)) \leftarrow V_{system\text{-}i}(S(t)) + \alpha_{i+}\delta(t+1) \quad \text{if } \delta(t+1) \geq 0, \text{ or}$$

$$V_{system\text{-}i}(S(t)) \leftarrow V_{system\text{-}i}(S(t)) + \alpha_{i-}\delta(t+1) \quad \text{if } \delta(t+1) < 0,$$

where α_{i+} and α_{i-} indicate the learning rates of the system i for positive (non-negative) and negative TD-RPEs, respectively, and they were systematically varied as described in the Results. The initial values of the system-specific state values of the system using the IR were set to 0.

In the system using the SR, the system-specific state value, $V_{system\text{-}i}(S(t))$, where i was 1 or 2, was calculated as a linear function of the feature variables of the state. In the SR, every state S_i was represented by a set of feature variables, which were estimated cumulative discounted future occupancies of all the states S_j (j = 1, ..., 25), denoted as $\sigma(S_i, S_j)$. These feature variables were updated through TD learning of SR features (Gershman et al., 2012; Gardner et al., 2018). Specifically, when the agent was at $S(t)$ at t and moved to $S(t+1)$ at $t+1$, where $S(t)$ was not the rewarded state, the TD

errors for SR features were calculated as:

$$\delta_{\mathrm{SR}j}(t+1) = 1 + \gamma\sigma(S(t+1), S_j) - \sigma(S(t), S_j) \quad \text{if } S(t) \text{ was } S_j, \text{ or}$$

$$\delta_{\mathrm{SR}j}(t+1) = 0 + \gamma\sigma(S(t+1), S_j) - \sigma(S(t), S_j) \quad \text{if } S(t) \text{ was not } S_j,$$

for all the states S_j (j = 1, ..., 25). If $S(t)$ was the rewarded state, the term $\gamma\sigma(S(t+1), S_j)$ was dropped. Using these TD errors, $\sigma(S(t), S_j)$ (j = 1, ..., 25) were updated as:

$$\sigma(S(t), S_j) \leftarrow \sigma(S(t), S_j) + \alpha_{\mathrm{SRfeature}}\delta_{\mathrm{SR}j}(t+1),$$

where $\alpha_{\mathrm{SRfeature}}$ indicates the learning rate for the update of SR features and was set to 0.05. The initial values of all the feature variables were set to 0. The system-specific state value, $V_{system\text{-}i}(S(t))$, was given as a linear function of the feature variables as:

$$V_{system\text{-}i}(S(t)) = \Sigma_j \{w_j\sigma(S(t), S_j)\}.$$

The weights (coefficients) w_j (j = 1, ..., 25) were updated based on the TD-RPE $\delta(t+1)$ described above, with potentially different learning rates depending on whether the TD-RPE was positive (non-negative) or negative:

$$w_j \leftarrow w_j + \alpha_{i+}\sigma(S(t), S_j)\delta(t+1) \quad \text{if } \delta(t+1) \geq 0, \text{ or}$$

$$w_j \leftarrow w_j + \alpha_{i-}\sigma(S(t), S_j)\delta(t+1) \quad \text{if } \delta(t+1) < 0,$$

where α_{i+} and α_{i-} indicate the learning rates of the system i for positive (non-negative) and negative TD-RPEs, respectively, and they were systematically varied as described in the Results. The initial values of the weights w_j were set to 0.

Simulations

Simulations were conducted 100 times for each condition by using MATLAB. Probabilities were implemented by "rand" and "randperm" functions of MATLAB.

Code availability

Codes to generate the data presented in the figures are available at:
https://github.com/kenjimoritagithub/sr

Results

Performance of combined SR-based and IR-based systems, with learning rate ratio varied

We first examined the case in which one system employed the SR whereas the other adopted the IR, systematically varying the ratio of learning rates from positive and negative TD-RPEs for each system, denoted as α_{SR+} / α_{SR-} for the SR-based system and α_{IR+} / α_{IR-} for the IR-based system. The sums of the learning rates from positive and negative TD-RPEs (i.e., $\alpha_{SR+} + \alpha_{SR-}$ and $\alpha_{IR+} + \alpha_{IR-}$) were kept constant at 1, and the inverse temperature and the time discount factor were also kept constant at $\beta = 5$ and $\gamma = 0.7$. We simulated the agent's behavior in a reward navigation task in a certain dynamic reward environment (Figure 1). We counted how many rewards the agent obtained during the task as a measure of performance, and examined how the obtained rewards, averaged across 100 simulations for each condition, varied depending on the learning rate ratios in both systems.

Figure 3 shows the results. As shown in the figure, the performance greatly varied depending on the conditions, and the best performance was achieved in the conditions in which α_{SR+} / α_{SR-} was high (much larger than 1) and α_{IR+} / α_{IR-} was low (much smaller than 1), i.e., the SR-based system learned mainly from positive TD-RPEs whereas the IR-based system learned mainly from negative TD-RPEs. We also examined the cases where the sums of the learning rates from positive and negative TD-RPEs were increased (1.25) or decreased (0.75), and also the inverse temperature and the time discount factor were varied ($\beta = 5$ or 10 and $\gamma = 0.7$ or 0.8). As shown in Figure 4, the result that combination of the SR-based system learning mainly from positive TD-RPEs and the IR-based system learning mainly from negative TD-RPEs achieved the best performance was preserved across these parameter changes.

Learning profiles of the model with different combinations of the learning rate ratios

Going back to the case with original parameter values used in Figure 3 ($\alpha_+ + \alpha_- = 1$, $\beta = 5$, and $\gamma = 0.7$), we examined the learning curves of the model in the case where the ratio of positive versus negative TD-RPE-based learning rate was respectively high (4) and low (0.25) in the SR-based and IR-based systems (i.e., the case achieving the best performance), compared with the cases where the ratio was 1 in both systems or the ratio was low (0.25) and high (4) in the SR-based and IR-based systems. Figure 5A shows the mean learning curves in each case, i.e., the mean time (number of time steps) used for obtaining the first, second, third, .. reward *placed* in each rewarded epoch, averaged across eight (from the second to the ninth) rewarded epochs and also across simulations (if a reward was placed in the first epoch and obtained in the second epoch, for example, it was regarded as the last reward in the first epoch rather than the first reward in the second epoch). Only the cases in which reward was obtained in not smaller than a quarter of (25 out of total 100) simulations were plotted, in

order to see the general average behavior of the model. Comparing the case with high α_{SR+} / α_{SR-} & low α_{IR+} / α_{IR-} (brown line) and the case with low α_{SR+} / α_{SR-} & high α_{IR+} / α_{IR-} (blue-green line), the time consumed for obtaining the first reward in an epoch, as well as the asymptotic value of the time spent for reward acquisition, were much longer in the latter case. The case with equal ratios for both systems (red line) showed an intermediate behavior.

We also looked at how each system developed system-specific values of the states, again in the three cases where α_+ / α_- for the SR-based and IR-based systems were 4 and 0.25, 1 and 1, or 0.25 and 4. Specifically, we looked at single-simulation examples of those values just after the agent obtained the last reward in the last (9th) rewarded epoch in single simulations, in which that reward was placed at the "special candidate state (with high probability of becoming a rewarded state: see the Methods)" in that epoch, in the three cases. Figure 5Ba shows the best-performance case where α_{SR+} / α_{SR-} was high (4) and α_{IR+} / α_{IR-} was low (0.25). The left panels show both the system-specific and integrated values, and the middle and right panels show only the integrated values in different ways. As shown in the figure, the system-specific state values of the SR-based system (red line in Figure 5Ba left) were all positive whereas those of the IR-based system (blue line) were all negative, and these two look largely symmetric. Meanwhile, the average of these two system-specific values, i.e., the integrated state values (black lines in Figure 5Ba left and middle; also shown as a heat map in Figure 5Ba right), have much smaller magnitudes but show a clear pattern, which is expected to approximate the true state values under the policy that the agent was taking. Figure 5Bb shows the case where both α_{SR+} / α_{SR-} and α_{IR+} / α_{IR-} were 1. In this case, both systems developed similar values. Figure 5Bc shows the case where α_{SR+} / α_{SR-} was low (0.25) and α_{IR+} / α_{IR-} was high (4). In this case, the SR- and IR-system-specific values were negative and positive, respectively, on the contrary to the case of Figure 5Ba with high α_{SR+} / α_{SR-} and low α_{IR+} / α_{IR-}. The integrated values have comparable magnitudes to the other two cases, but were negative in many states. Comparing the right panels of Figure 5Ba and Figure 5Bc, the peak of the value function looks sharper in the latter, presumably reflecting the smaller degree of positive TD-RPE-dependent learning in the SR-based system.

Cases where both systems employ only the IR or only the SR

We also examined the cases where both systems employed only the SR, or only the IR. Figure 6A and 6B show the performance results for the SR only and IR only cases, respectively, for the same sets of parameters as shown in Figure 4 for the case of SR and IR combination. Notably, if the ratio of positive versus negative TD-RPE-based learning rates was equal between both systems that employ only the SR or IR, the two systems behaved in exactly the same manner, and thus such conditions (on the vertical lines in Figure 6A and 6B) were equivalent to having only a single system. Also, Figure 6A and 6B only show either the left or right side, because "α_{1+} / α_{1-} = 0.2 and α_{2+} / α_{2-} = 3", for example, are equivalent to "α_{1+} / α_{1-} = 3 and α_{2+} / α_{2-} = 0.2" given that both systems 1 and 2 employed the same representation (SR or IR) so that we examined only one of these, and Figure 6A

shows the left side whereas Figure 6B shows the right side just because such arrangements might facilitate visual comparisons with Figure 4, where the ratios for the SR-based system and the IR-based system were plotted on the axes rising to the left and the right, respectively.

Comparing Figure 6A and 6B, SR-SR combination generally achieved better performance than IR-IR combination. This may not be surprising, given that a known feature of SR-based learning is sensitive adaptation to changes in rewards in the environments (Momennejad et al., 2017; Russek et al., 2017), which occurred in our simulated task. Besides, during the initial no-reward epoch, the SR-based system could acquire, through TD learning of SR features, an SR under the random policy, which presumably acted as beneficial "latent learning" (Dayan, 1993), while the IR-based system could not learn anything. Likewise, although the location of special reward candidate state and thus the optimal policy varied from epoch to epoch and SR is a policy-dependent representation, SR features acquired and updated in the agent should contain information about the basic structure of the grid space, which could help reward navigation throughout the task. Next, whether having two systems with different ratios of positive and negative TD-RPE-based learning rates was advantageous appears to be less clear in the cases where both systems employed only the SR or IR than in the case of the combination of SR- and IR-based systems. However, the best-performance ratio set (white cross in the figure) was actually not on the vertical line in most conditions, suggesting that having two systems with different ratios was in fact beneficial. Last and most strikingly, the best combination of the SR-based and IR-based systems, i.e., the combination of the SR-based system learning mainly from positive TD-RPEs and the IR-based system learning mainly from negative TD-RPEs outperformed any combinations of two SR-based systems or two IR-based systems for all the examined sets of parameters ($\alpha_+ + \alpha_-$, β, and γ) (compare Figure 4 and Figure 6A,B; also shown in Figure 6C).

Dependence on task properties

So far we examined the agent's performance in the particular task. We examined how it could differ if task properties change. The original task contained a stochastic nature in a sense that reward was placed at a particular state (named the special reward candidate state) with a high probability (60%) in a given epoch but with the remaining probability reward was placed at one of the other (normal) candidate states. We examined what happened when this stochasticity was reduced or removed in the case of the model consisting of SR-based and IR-based systems with the original set of parameters used in Figures 3 and 5 ($\alpha_{SR+} + \alpha_{SR-} = \alpha_{IR+} + \alpha_{IR-} = 1$, $\beta = 5$, and $\gamma = 0.7$). Figure 7A shows the results for the cases where the probability of reward placement at the special candidate state was varied from 70% to 100%. As shown in the figure, as the stochasticity reduced, combinations of α_{SR+} / α_{SR-} and α_{IR+} / α_{IR-} that gave good performance gradually shifted, and when the stochasticity was totally removed, best performance was achieved when both α_{SR+} / α_{SR-} and α_{IR+} / α_{IR-} were high values and actually α_{IR+} / α_{IR-} was even higher than α_{SR+} / α_{SR-}.

We also examined the effects of another variation in the task properties. In the original task, new reward was not placed until the agent obtained the previous reward. This property, coupled with the stochastic placement of reward, is considered to make perseveration-like behavior quite maladaptive and increase the importance of learning from negative feedbacks. We examined what happened if this property was weakened by introducing periodic resets of reward placement into the original task, again in the case of the model consisting of SR-based and IR-based systems with the original set of parameters. Figure 7B shows the results for the cases where reward location (state) was reset at every 500, 250, 100, or 50 time steps in the rewarded epochs; at the reset timings, reward was placed at the special candidate state for the epoch with 60% and placed at one of the other (normal) candidate states with 5% each. As shown in the figure, as the resets became more frequent, combinations of α_{SR+} / α_{SR-} and α_{IR+} / α_{IR-} that gave good performance again gradually shifted from high α_{SR+} / α_{SR-} and low α_{IR+} / $\alpha_{IR-}$to high α_{SR+} / α_{SR-} and also high α_{IR+} / α_{IR-}.

We further examined three different variations of the original task. The first variation was abolishment of multiple rewarded epochs with different placements of special reward candidate state. Specifically, we examined the case where there was only a single rewarded epoch with 4500 time steps, instead of nine epochs with 500 time steps for each, following the initial 500 time-steps no-reward epoch. The special candidate state was varied only across simulations. The second variation was removal of the stochasticity of reward placement within each epoch, together with shortening of the duration of each rewarded epoch (from 500 time steps to 100 time steps) and increasing the number of rewarded epochs (from 9 to 45). The third variation was abolishment of special reward candidate state. Specifically, we came back to the original nine 500 time-steps rewarded epochs, but the probability of reward placement at the special candidate state was changed from the original 60% to 1/9 (11.11..%) so that reward was placed at each of the nine candidate states with equal probability (1/9). Figures 7C-E show the performance of the model consisting of SR-based and IR-based systems with the original set of parameters in these three variations. As shown in the figures, the patterns look similar to the case of the original task (Figure 3).

It was rather surprising that similar pattern appeared even in the last case, because reward was placed at any candidate states with equal probabilities and so it was no longer a task where the agent could learn where was the special candidate state in each epoch. However, the nine candidate states were neighboring with each other on the two edges of the grid space, and thus it would still have been the case that next reward was placed at a state near from the previously rewarded state with a relatively high probability and the agent could exploit it through positive TD-RPE-based learning. As a control, we examined a task in which positive TD-RPE-based learning was expected to have little meaning. Specifically, after 500 time-steps no-reward epoch, rewarded state was determined by a fixed order, namely, (5, 1), (5, 5), (1, 5), and again (5, 1), and these were repeated throughout the task. Figure 7F shows the performance of the model consisting of SR-based and IR-based systems with the original parameters. As shown in the figure, best performance, though rather low as expected, was achieved

when both α_{SR+} / α_{SR-} and α_{IR+} / α_{IR-} were low values. This is reasonable because only quick erasing of the memory of previous reward is considered to be adaptive in this control task.

These results together indicate that the combination of appetitive SR and aversive IR systems does not always achieve the best performance, but does so in certain environments where reward placement is dynamically changing but still learnable using both positive and negative TD-RPEs. The reason for this will be discussed at the beginning of the Discussion.

Discussion

We found that combination of SR-based system learning mainly from positive TD-RPEs and IR-based system learning mainly from negative TD-RPEs showed superior performance in certain situations where reward placement was dynamically changing. The reason why such a combination achieved the best performance is considered to be as follows. When the agent unexpectedly encountered a reward at a certain state, next reward was placed at the same state or nearby states with a certain high probability (given that there was learnable structure). Then, quickly revising up the value of not only the rewarded state itself but also any states from which the rewarded state could be easily reached (i.e., states having high SR feature values) would be beneficial. In other words, SR-dependent generalization of positive TD-RPE-based value updates would be beneficial. Next, assume that the agent obtained rewards several times at the same or nearby states, and then reward was placed at a distant state. This time, the agent was likely to unexpectedly encounter reward omission. Critically, through the repetitive acquisition of rewards at the nearby states, the agent's policy was likely to be sharpened and therefore revising down the values of only the states right on the sharpened policy would already be effective. In other words, as for negative TD-RPE-based value updates, IR-based narrow effect would be largely sufficient, and too much SR-dependent generalization could rather be harmful.

Appetitive and aversive learning in distinct neural pathways

There have been a number of empirical studies suggesting that the D1/direct and D2/indirect pathways of BG are primarily involved in appetitive and aversive learning, respectively. Parkinson disease patients on and off medication showed opposite patterns of learning, better from positive and negative feedbacks, respectively, presumably through changes in the balance of D1/direct and D2/indirect pathways (Frank et al., 2004). Optogenetic stimulation of D1 or D2 SPNs induced reinforcement or punishment (Kravitz et al., 2012), and mimicked increase or decrease in action values (Tai et al., 2012), respectively. Blockade of the direct or indirect pathway impaired reward or avoidance learning, respectively (Hikida et al., 2010). At the synaptic level, D1R- and D2R-mediated plasticities have opposite DA dependences (Shen et al., 2008), potentiated by phasic DA rise (Yagishita et al., 2014) and dip (Iino et al., 2020), respectively.

Theoretical studies explored potential functional merits of having such distinct BG pathways for appetitive and aversive learning. A study (Collins and Frank, 2014) suggested that the D1/direct and D2/indirect pathways could learn the benefits and costs of the same action, respectively, and DA-dependent differential weighting of costs and benefits could be achieved. Other studies (Mikhael and Bogacz, 2016) suggested that these pathways could learn not only the mean but also the uncertainty of reward, and the set of plasticity rules achieving it could also realize estimation of costs and payoffs (Möller and Bogacz, 2019).

Differences in the inputs to the neural pathways for appetitive and aversive learning

These existing models assumed that the D1/direct and D2/indirect pathways employ the same way of representation of states or actions, specifically, the IR. Since different neural regions or populations may generally adopt different ways of representation, as in the case of egocentric and allocentric spatial representations (Chen et al., 2014; Town et al., 2017; Wang et al., 2020), the two BG pathways might also use different representations if their input sources are not identical.

There are two major types of cortical pyramidal cells projecting to the striatum, namely, intratelencephalic (IT) neurons, which project to bilateral striatum, and pyramidal-tract (PT) neurons, which target ipsilateral striatum (Cowan and Wilson, 1994; Reiner et al., 2010). Electrophysiological and optogenetic studies have shown that both D1/direct and D2/indirect pathways receive inputs from both IT and PT neurons (Ballion et al., 2008; Kress et al., 2013). However, anatomical studies suggested that D1 and D2 SPNs predominantly receive inputs from IT and PT neurons, respectively (Lei et al., 2004; Reiner et al., 2010). Model fitting of previously published paired-pulse ratio data also suggested IT→D1 and PT→D2 preferential activation through short-term synaptic plasticity (Morita, 2014). At a larger spatial scale, it has been shown that different cortical regions innervate the D1/direct and D2/indirect pathways in different proportions (Wall et al., 2013). In particular, the somatosensory and limbic cortical regions preferentially innervated the direct pathway whereas the primary motor cortex preferentially targeted the indirect pathway.

Given these, it is conceivable that the D1/direct and D2/indirect pathways unevenly receive inputs from different sources, which could represent states in different ways. Other than the BG pathways, the amygdala-OFC and OFC-NAc pathways have also been suggested to be specifically involved in appetitive and aversive learning, respectively (Groman et al., 2019). In this case, the pathways have different input structures (i.e., amygdala and OFC), which may use different state representations.

What types of state representation are used in the different input neuron types or regions?

Given that the D1/direct and D2/indirect pathways are unevenly innervated by the IT and PT cortical neurons and/or by the sensory/limbic and motor cortices as suggested, an emerging question is whether either type of state representation, in particular, SR or IR, is preferentially used in the IT and PT neurons or in the sensory/limbic and motor cortices.

As for neuron types, recent studies have suggested that sensitivity to outcome devaluation, which is indicative of model-based/goal-directed behavior, is mediated via IT, rather than PT, neurons projecting from the prelimbic cortex to the posterior dorsomedial striatum (Hart et al., 2018a; Hart et al., 2018b). Since SR-based TD learning also enables sensitive adaptation to changes in outcome

values (Momennejad et al., 2017; Russek et al., 2017), these results would also be consistent with the possibility that IT neurons, but not PT neurons, are engaged in SR-like representation, while the results also suggest that SR-based learning should occur from negative feedbacks (i.e., devaluation) at least to some extent. Local connections between IT neurons are less reciprocal than those between PT neurons (Morishima et al., 2011), and axonal projections of individual IT neurons to the striatum tend to have more widespread endpoints than those of PT neurons (Morita et al., 2019). These properties might also be in line with the possibility that IT neurons have SR-like, rather than punctate IR-like, state representation. Regarding cortical regions, BOLD signals consistent with SR have been observed in hippocampal–entorhinal cortex (Garvert et al., 2017), and recent work found evidence for SR in decoded activity in sensory cortex (Russek et al., 2021).

Given that IT neurons and/or sensory/limbic cortices preferentially innervate the BG direct pathway while PT neurons and/or motor cortex preferentially innervate the indirect pathway, the suggested possible employment of SR in IT neurons or in the hippocampal–entorhinal or sensory cortices indicates that the direct pathway preferentially employs the SR while the indirect pathway does not (Figure 8). Intriguingly, this matches the best-performance combination in our simulations: SR-based and IR-based systems learning mainly from positive and negative TD-RPEs, respectively. It is tempting to speculate that such a combination has been equipped in the brain through evolution because it is adaptive. In the best combination, the SR-based system-specific values became positive while the IR-based system specific values became negative, and these two looked largely symmetric (Figure 5Ba). If these systems are implemented by the BG direct and indirect pathways, respectively, such symmetry means that these pathways show largely similar activity patterns. This seems in line with the observation (Cui et al., 2013) that the striatal direct- and indirect-pathway neurons showed concurrent activation, rather than antagonistic activation that was expected from the classical view.

Strengths, limitations, and perspectives

As discussed above, our dual systems model with the best-performance combination can potentially be well mapped onto the neuron type- and/or cortical area-specific circuitry in the cortico-BG system, going beyond the previous models that took into account the two BG pathways but not the diversity and complexity of the neocortex. Our model is still algorithmic, but it could be a step onto the anticipated bridge from algorithmic to implementational levels (Averbeck and O'Doherty, 2021).

In the present work, we assumed that TD-RPEs are calculated from the integrated values, but how it can be achieved remains elusive. Previously two of us proposed that the IT→D1/direct and PT→D2/indirect pathways respectively represent the current and previous values and they are used for TD-RPE calculation (Morita et al., 2012; Morita et al., 2013), but we considered different possibilities about these pathways in the present work. Integrating our model with mechanisms of TD error calculation (Watabe-Uchida et al., 2017) is an important future direction.

Our results suggest that combination of SR-based learning mainly from positive TD-RPEs and IR-based learning mainly from negative TD-RPEs is a good learning strategy. Whether humans or animals indeed take such a strategy can potentially be examined through behavioral experiments to examine whether the behavior exhibits features of SR-based learning (c.f., (Momennejad et al., 2017)) differently between the cases with positive and negative TD-RPEs. Brain substrates can also be explored by neuroimaging (c.f., (Garvert et al., 2017; Russek et al., 2021)).

References

Averbeck B, O'Doherty JP (2021) Reinforcement-learning in fronto-striatal circuits. Neuropsychopharmacology:10.1038/s41386-41021-01108-41380.

Balleine BW, O'Doherty JP (2010) Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. Neuropsychopharmacology 35:48-69.

Ballion B, Mallet N, Bézard E, Lanciego JL, Gonon F (2008) Intratelencephalic corticostriatal neurons equally excite striatonigral and striatopallidal neurons and their discharge activity is selectively reduced in experimental parkinsonism. Eur J Neurosci 27:2313-2321.

Chen Y, Monaco S, Byrne P, Yan X, Henriques DY, Crawford JD (2014) Allocentric versus egocentric representation of remembered reach targets in human cortex. J Neurosci 34:12515-12526.

Collins AG, Frank MJ (2014) Opponent actor learning (OpAL): modeling interactive effects of striatal dopamine on reinforcement learning and choice incentive. Psychol Rev 121:337-366.

Cowan RL, Wilson CJ (1994) Spontaneous firing patterns and axonal projections of single corticostriatal neurons in the rat medial agranular cortex. J Neurophysiol 71:17-32.

Cui G, Jun SB, Jin X, Pham MD, Vogel SS, Lovinger DM, Costa RM (2013) Concurrent activation of striatal direct and indirect pathways during action initiation. Nature 494:238-242.

Dayan P (1993) Improving Generalization for Temporal Difference Learning: The Successor Representation. Neural Computation 5:613-624.

Dolan RJ, Dayan P (2013) Goals and habits in the brain. Neuron 80:312-325.

Frank MJ, Seeberger LC, O'reilly RC (2004) By carrot or by stick: cognitive reinforcement learning in parkinsonism. Science 306:1940-1943.

Gardner MPH, Schoenbaum G, Gershman SJ (2018) Rethinking dopamine as generalized prediction error. Proc Biol Sci 285:20181645.

Garvert MM, Dolan RJ, Behrens TE (2017) A map of abstract relational knowledge in the human hippocampal-entorhinal cortex. Elife 6:e17086.

Gershman SJ, Moore CD, Todd MT, Norman KA, Sederberg PB (2012) The successor representation and temporal context. Neural Comput 24:1553-1568.

Groman SM, Keistler C, Keip AJ, Hammarlund E, DiLeone RJ, Pittenger C, Lee D, Taylor JR (2019) Orbitofrontal Circuits Control Multiple Reinforcement-Learning Processes. Neuron 103:734-746.e733.

Hart G, Bradfield LA, Balleine BW (2018a) Prefrontal Corticostriatal Disconnection Blocks the Acquisition of Goal-Directed Action. J Neurosci 38:1311-1322.

Hart G, Bradfield LA, Fok SY, Chieng B, Balleine BW (2018b) The Bilateral Prefronto-striatal Pathway Is Necessary for Learning New Goal-Directed Actions. Curr Biol 28:2218-2229.e2217.

Hikida T, Kimura K, Wada N, Funabiki K, Nakanishi S (2010) Distinct roles of synaptic transmission

in direct and indirect striatal pathways to reward and aversive behavior. Neuron 66:896-907.

Iino Y, Sawada T, Yamaguchi K, Tajiri M, Ishii S, Kasai H, Yagishita S (2020) Dopamine D2 receptors in discrimination learning and spine enlargement. Nature 579:555-560.

Kravitz AV, Tye LD, Kreitzer AC (2012) Distinct roles for direct and indirect pathway striatal neurons in reinforcement. Nat Neurosci 15:816-818.

Kress GJ, Yamawaki N, Wokosin DL, Wickersham IR, Shepherd GM, Surmeier DJ (2013) Convergent cortical innervation of striatal projection neurons. Nat Neurosci 16:665-667.

Lei W, Jiao Y, Del Mar N, Reiner A (2004) Evidence for differential cortical input to direct pathway versus indirect pathway striatal projection neurons in rats. J Neurosci 24:8289-8299.

Mikhael JG, Bogacz R (2016) Learning Reward Uncertainty in the Basal Ganglia. PLoS Comput Biol 12:e1005062.

Momennejad I, Russek EM, Cheong JH, Botvinick MM, Daw ND, Gershman SJ (2017) The successor representation in human reinforcement learning. Nat Hum Behav 1:680-692.

Morishima M, Morita K, Kubota Y, Kawaguchi Y (2011) Highly differentiated projection-specific cortical subnetworks. J Neurosci 31:10380-10391.

Morita K (2014) Differential cortical activation of the striatal direct and indirect pathway cells: reconciling the anatomical and optogenetic results by using a computational method. J Neurophysiol 112:120-146.

Morita K, Im S, Kawaguchi Y (2019) Differential Striatal Axonal Arborizations of the Intratelencephalic and Pyramidal-Tract Neurons: Analysis of the Data in the MouseLight Database. Front Neural Circuits 13:71.

Morita K, Morishima M, Sakai K, Kawaguchi Y (2012) Reinforcement learning: computing the temporal difference of values via distinct corticostriatal pathways. Trends Neurosci 35:457–467.

Morita K, Morishima M, Sakai K, Kawaguchi Y (2013) Dopaminergic control of motivation and reinforcement learning: a closed-circuit account for reward-oriented behavior. J Neurosci 33:8866-8890.

Möller M, Bogacz R (2019) Learning the payoffs and costs of actions. PLoS Comput Biol 15:e1006285.

Reiner A, Hart NM, Lei W, Deng Y (2010) Corticostriatal projection neurons - dichotomous types and dichotomous functions. Front Neuroanat 4:142.

Russek EM, Momennejad I, Botvinick MM, Gershman SJ, Daw ND (2017) Predictive representations can link model-based reinforcement learning to model-free mechanisms. PLoS Comput Biol 13:e1005768.

Russek EM, Momennejad I, Botvinick MM, Gershman SJ, Daw ND (2021) Neural evidence for the successor representation in choice evaluation. bioRxiv https://doi.org/10.1101/2021.08.29.458114.

Shen W, Flajolet M, Greengard P, Surmeier DJ (2008) Dichotomous dopaminergic control of striatal synaptic plasticity. Science 321:848-851.

Stachenfeld KL, Botvinick MM, Gershman SJ (2017) The hippocampus as a predictive map. Nat Neurosci 20:1643-1653.

Sutton R, Barto A (1998) Reinforcement Learning. Cambridge, MA: MIT Press.

Tai LH, Lee AM, Benavidez N, Bonci A, Wilbrecht L (2012) Transient stimulation of distinct subpopulations of striatal neurons mimics changes in action value. Nat Neurosci 15:1281-1289.

Town SM, Brimijoin WO, Bizley JK (2017) Egocentric and allocentric representations in auditory cortex. PLoS Biol 15:e2001878.

Wall NR, De La Parra M, Callaway EM, Kreitzer AC (2013) Differential innervation of direct- and indirect-pathway striatal projection neurons. Neuron 79:347-360.

Wang C, Chen X, Knierim JJ (2020) Egocentric and allocentric representations of space in the rodent brain. Curr Opin Neurobiol 60:12-20.

Watabe-Uchida M, Eshel N, Uchida N (2017) Neural Circuitry of Reward Prediction Error. Annu Rev Neurosci 40:373-394.

Yagishita S, Hayashi-Takagi A, Ellis-Davies GC, Urakubo H, Ishii S, Kasai H (2014) A critical time window for dopamine actions on the structural plasticity of dendritic spines. Science 345:1616-1620.

Figure legends

Figure 1

The simulated reward navigation task. **(A)** The 5 × 5 grid space, where the agent moved around. The agent started from the fixed start state, (1, 1), and moved to one of the neighboring states (two states at the corners, three states at the edges, and four states elsewhere) at each time step. There were nine reward candidate states, where reward could potentially be placed, namely, (1, 5), (2, 5), (3, 5), (4, 5), (5, 1), (5, 2), (5, 3), (5, 4), and (5, 5) (indicated by the gray color). **(B)** Epochs in the task. During the initial 500 time steps, there was no reward, and this was called the no-reward epoch. During the next 500 time steps, one of the nine reward candidate states was specified as a special candidate state, whereas the remaining eight reward candidate states were regarded as normal candidate states. There were in total nine rewarded epochs (500 time steps for each), and each one of the nine reward candidate states became the special candidate state in one of the nine epochs; the order was determined by pseudorandom permutation in each single simulation. **(C)** In the rewarded epochs, if the agent reached the rewarded state and obtained the reward, the agent was immediately carried back to the start state, and a new reward was introduced into a state, which was the special reward candidate state with 60% probability and one of the eight normal candidate states with 5% probability for each.

Figure 2

The reinforcement learning model with two systems. **(A)** Model architecture. The model consists of two learning systems, system 1 and system 2, which may use different ways of state representations. Each system has its own system-specific value of each state, and their mean (average) becomes the integrated state value. The agent selects an action to move to a neighboring state depending on their integrated values in a soft-max manner. The integrated values are also used for calculating the temporal-difference reward prediction error (TD-RPE). The TD-RPE was then used for updating the system-specific state values in each of the two systems. The leaning rates of the updates of the system-specific state values can differ between the two systems and also depending on whether the TD-RPE is positive (non-negative) or negative. **(B)** Combinations of the ways of state representation in the two systems. Combination of the successor representation (SR) in system 1 and the individual representation (IR) in system 2 was examined in Figures 3-5 and 7. Cases with only the SR or IR in both systems were examined in Figure 6.

Figure 3

Performance of the model consisting of a system using the SR and another system using the IR, and its dependence on the ratios of the learning rates from positive and negative TD-RPEs in each system. The axis rising to the left indicates the ratio of positive- / negative-error-based learning rates (denoted as α_+/α_-) in the system using the SR, while the axis rising to the right indicates the same ratio in the

system using the IR. The sum of the learning rates from positive and negative TD-RPEs ($\alpha_+ + \alpha_-$) in each system was constant at 1 in any conditions shown in this figure. The inverse temperature β was 5, and the time discount factor γ was 0.7. The vertical line corresponds to conditions where the α_+/α_- ratio is equal for both systems (bottom: negative error-based learning dominates; top: positive error-based learning dominates). The left side of the vertical line corresponds to conditions where the α_+/α_- ratio is larger in the SR-based system than in the IR-based system, whereas the opposite is the case for the right side of the vertical line. The color in each pixel indicates the total obtained rewards in the task, averaged across 100 simulations for each condition, in reference to the rightmost color bar. The white cross indicates the set of α_+/α_- ratios that gave the best performance ($\alpha_+/\alpha_- = 4$ and 0.25 for the SR- and IR-based systems, respectively) among those examined in this figure. Note that the minimum value of the color bar is not 0. Also, the maximum and minimum values of the color bar do not match the highest and lowest performances in this figure; rather, they were set so that the highest and lowest performances in all the simulations of the same task with different parameters and/or model architecture (i.e., not only SR+IR but also SR+SR and IR+IR) shown in this figure and Figures 4 and 6 can be covered.

Figure 4

Performance of the model consisting of an SR-based system and an IR-based system with various sets of parameters. The sum of the learning rates from positive and negative TD-RPEs ($\alpha_+ + \alpha_-$) was varied over 0.75, 1, and 1.25. The inverse temperature β was varied over 5 and 10. The time discount factor γ was varied over 0.7 and 0.8. The ratio of positive- / negative-error-based learning rates (α_+/α_-) was varied over 0.2, 0.25, 1/3, 0.5, 1, 2, 3, 4, and 5 for the cases with $\alpha_+ + \alpha_- = 0.75$ or 1, but 0.2 and 5 were omitted for the cases with $\alpha_+ + \alpha_- = 1.25$ in order to avoid learning rate larger than 1. The color-performance correspondence is the same as in Figure 3. The white cross in each panel indicates the set of α_+/α_- ratios that gave the best performance among those examined in that condition, and the gray number near the top of each panel indicates the best performance (total rewards). The panel with $\alpha_+ + \alpha_- = 1$, $\beta = 5$, and $\gamma = 0.7$ shows the same results as shown in Figure 3.

Figure 5

Learning profiles of the model consisting of an SR-based system and an IR-based system with different combinations of the ratios of positive- / negative-error-based learning rates in each system. Three cases included in Figure 3 (where the sum of the learning rates from positive and negative TD-RPEs ($\alpha_+ + \alpha_-$) in each system was 1, the inverse temperature β was 5, and the time discount factor γ was 0.7) were analyzed. **(A)** Mean learning curves. The curves indicate the mean time (number of time steps) used for obtaining the first, second, third, .. reward placed in each rewarded epoch (horizontal axis), averaged across eight (from the second to the ninth) rewarded epochs and also across simulations. Only the cases in which reward was obtained in not smaller than a quarter of (25 out of total 100) simulations were plotted. The brown, red, and blue-green curves correspond to the conditions with

(α_+/α_- for SR-system, α_+/α_- for IR-system) = (4, 0.25), (1, 1), and (0.25, 4), respectively (the colors match those in the corresponding pixels in Figure 3). **(B)** Examples of the system-specific state values and the integrated state values. Panels in (a), (b), and (c) correspond to the conditions with (α_+/α_- for SR-system, α_+/α_- for IR-system) = (4, 0.25), (1, 1), and (0.25, 4), respectively. The left panels show the system-specific state values (red line: SR-based system, blue line: IR-based system) and the integrated state values (black solid line). The horizontal axis indicates 25 states (1-5 correspond to (1,1)-(5,1) in Figure 1/5B-right; 6-10 correspond to (1,2)-(5,2), and so on). Values just after the agent obtained the last reward in the last (9th) rewarded epoch (indicated by the white crosses in the right panels) in single simulations, in which that reward was placed at the "special candidate state (with high probability of becoming a rewarded state: see the Methods)" in that epoch, were shown. Note the differences in the vertical scales between (a) or (c) and (b). The middle panels show only the integrated state values. The vertical axes in (a) and (c) were enlarged as compared to the left panels. The right panels also show the integrated state values as a heat map on the 5 × 5 grid space.

Figure 6

Performance of the model consisting of two systems, both of which employed the same way of state representation. Results in the case where both systems employed only the SR **(A)** or only the IR **(B)** are shown. Configurations are the same as those in Figure 4 (white cross: the best-performance set of α_+/α_- ratios; gray number: the best performance (total rewards)). If the α_+/α_- ratio was equal between both systems, the two systems behaved in exactly the same manner, and thus such conditions (on the vertical lines) were equivalent to having only a single system. Also, only the left (A) or right (B) side is shown, because "$\alpha_{1+} / \alpha_{1-} = 0.2$ and $\alpha_{2+} / \alpha_{2-} = 3$", for example, are equivalent to "$\alpha_{1+} / \alpha_{1-} = 3$ and $\alpha_{2+} / \alpha_{2-} = 0.2$" given that both systems employed the same representation. The left (A)-and right-(B) placement was made so as to facilitate visual comparisons with Figure 3. **(C)** Comparison of the maximum performances (total rewards) of the cases where one system employed the SR while the other employed the IR (purple), both systems employed the SR (red), and both systems employed the IR (blue) for each set of the time discount factor (γ), inverse temperature (β), and the sum of learning rates for positive and negative TD-RPEs ($\alpha_+ + \alpha_-$).

Figure 7

Performance of the model consisting of SR-based and IR-based system when task properties were changed. The model with the original set of parameters ($\alpha_{SR+} + \alpha_{SR-} = \alpha_{IR+} + \alpha_{IR-} = 1$, $\beta = 5$, and $\gamma = 0.7$) was used. The ranges of the color bars in this figure correspond to the ranges between the lowest and highest performances (i.e., the minimum and maximum total rewards) in the individual panels, and they are different from the color bars in the previous figures. **(A)** Performance (total rewards) for the cases where the probability of reward placement at the special candidate state was varied from 70% (leftmost panel), 80%, 90%, or 100% (rightmost). **(B)** Performance for the cases where reward location (state) was reset at every 500 (leftmost panel), 250, 100, or 50 (rightmost) time steps in the rewarded

epochs; where reward was located was determined according to the original stochastic rule, i.e., reward was placed at the special candidate state for the epoch with 60% and placed at one of the other (normal) candidate states with 5% each. **(C)** Performance for the case where there was only a single rewarded epoch with 4500 time steps following the initial 500 time steps no-reward epoch. **(D)** Performance for the case where stochasticity of reward placement within each epoch was removed and the duration of each rewarded epoch was shortened from 500 time steps to 100 time steps while number of rewarded epochs was increased from 9 to 45. **(E)** Performance for the case where special candidate state was abolished and reward was placed at each of the nine candidate states with equal probability (1/9 = 11.11..%). **(F)** Performance for the case where rewarded state was determined by a fixed order, namely, (5, 1), (5, 5), (1, 5), and again (5, 1), and these were repeated throughout the task.

Figure 8

Possible implementation of the SR-based and IR-based learning systems in the cortico-BG circuits. SR is used in the sensory and limbic cortices and/or intratelencephalic (IT) neurons, whereas IR is used in the primary motor cortex and/or pyramidal-tract (PT) neurons. Sensory and limbic cortices and/or IT neurons preferentially project to the D1/direct pathway, whereas primary motor cortex and/or PT neurons preferentially target the D2/indirect pathway. The D1/direct pathway learns mainly from positive TD-RPEs, whereas the D2/indirect pathway learns mainly from negative TD-RPEs. Overall, the circuits implement the combination of SR-based system learning mainly from positive TD-RPEs and IR-based system learning mainly from negative TD-RPEs, which achieved the best performance in our simulated task.

(End)

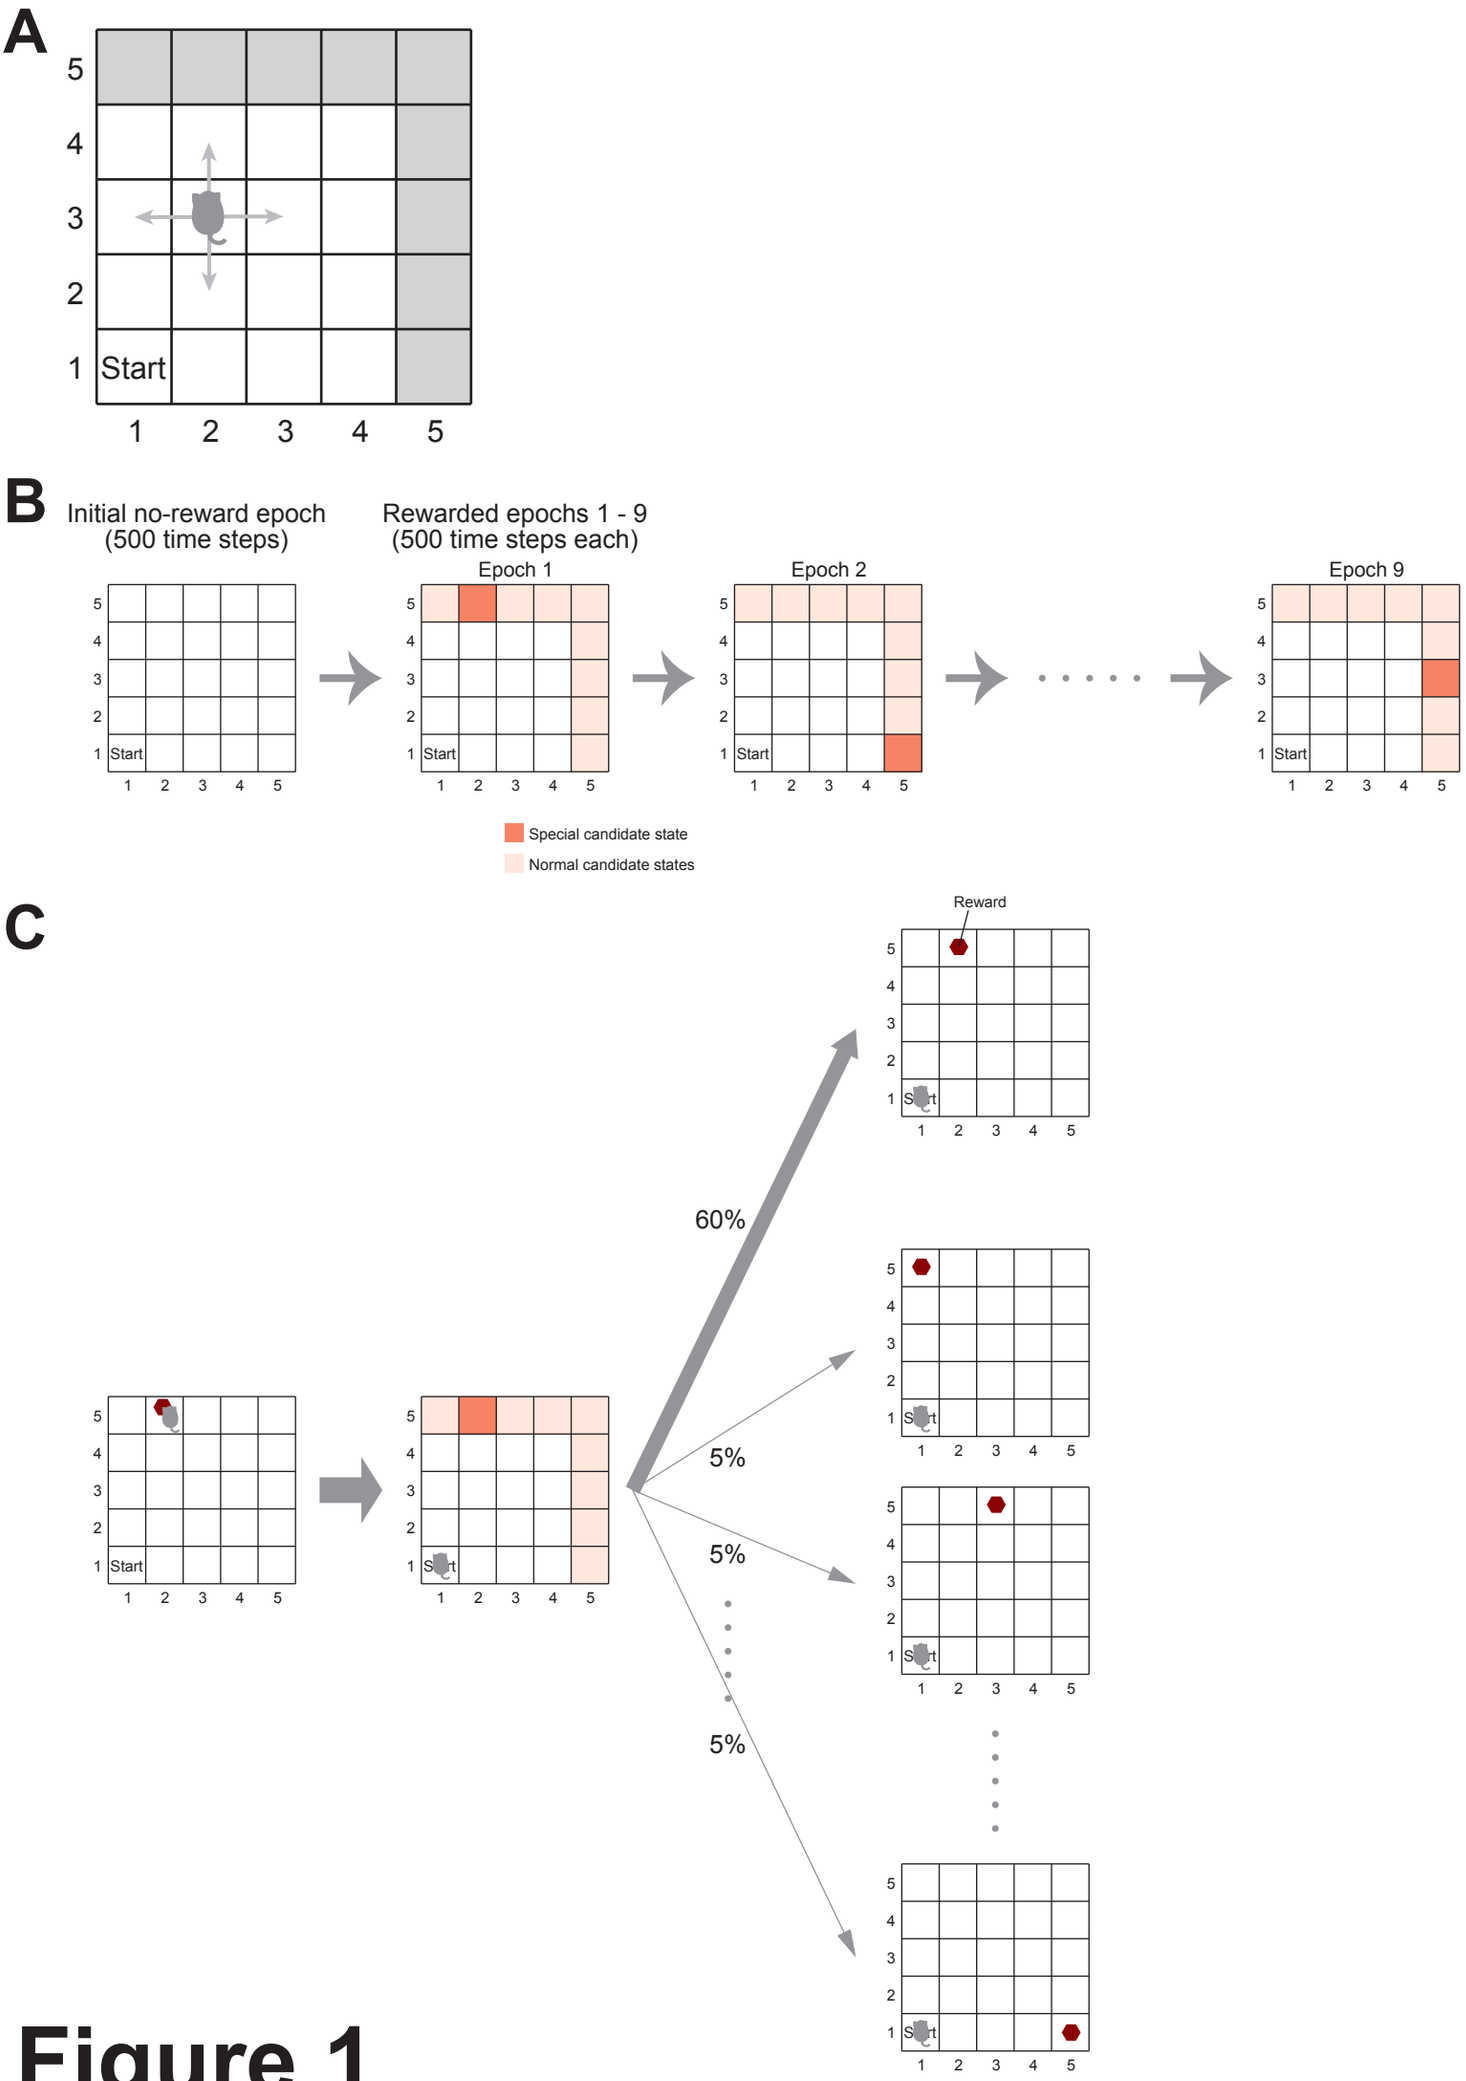
A
Start
B
Initial no-reward epoch
(500 time steps)
Rewarded epochs 1 - 9
(500 time steps each)
Epoch 1
Epoch 2
Epoch 9
Special candidate state
Normal candidate states
C
Reward
60%
5%
5%
5%

Figure 1

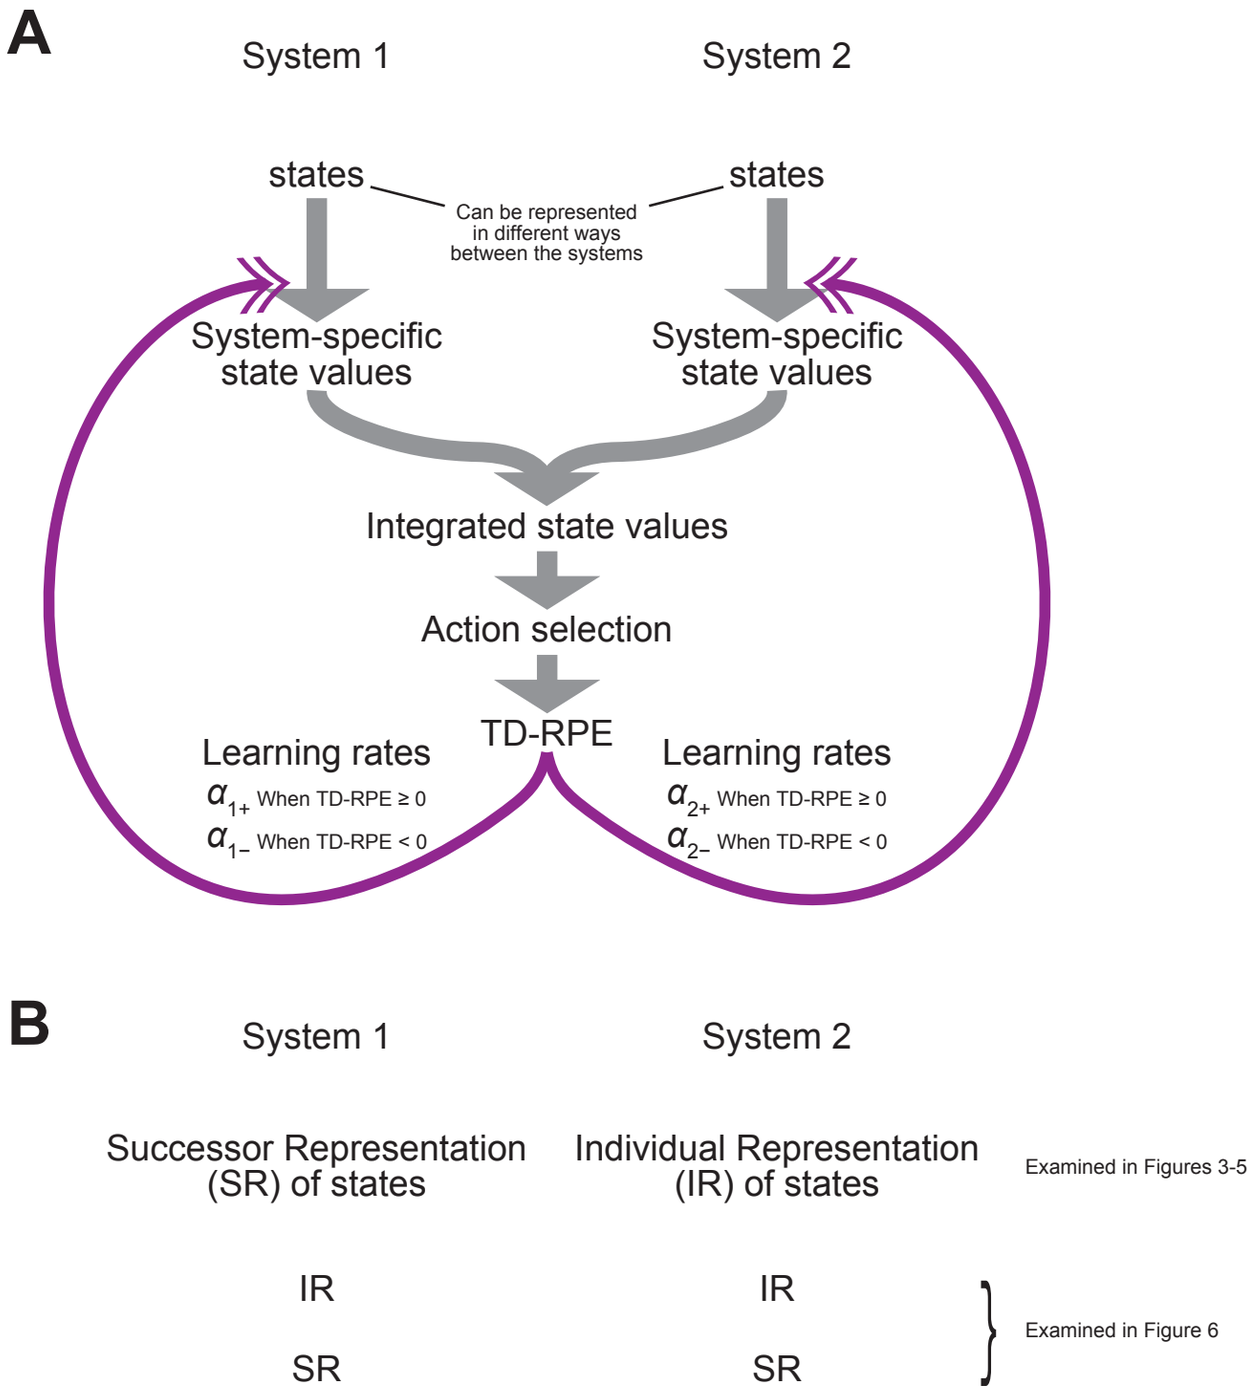
A
System 1
System 2
states
states
Can be represented in different ways between the systems
System-specific state values
System-specific state values
Integrated state values
Action selection
TD-RPE
Learning rates
α_{1+} When TD-RPE ≥ 0
α_{1-} When TD-RPE < 0
Learning rates
α_{2+} When TD-RPE ≥ 0
α_{2-} When TD-RPE < 0
B
System 1
System 2
Successor Representation (SR) of states
Individual Representation (IR) of states
Examined in Figures 3-5
IR
IR
SR
SR
Examined in Figure 6

Figure 2

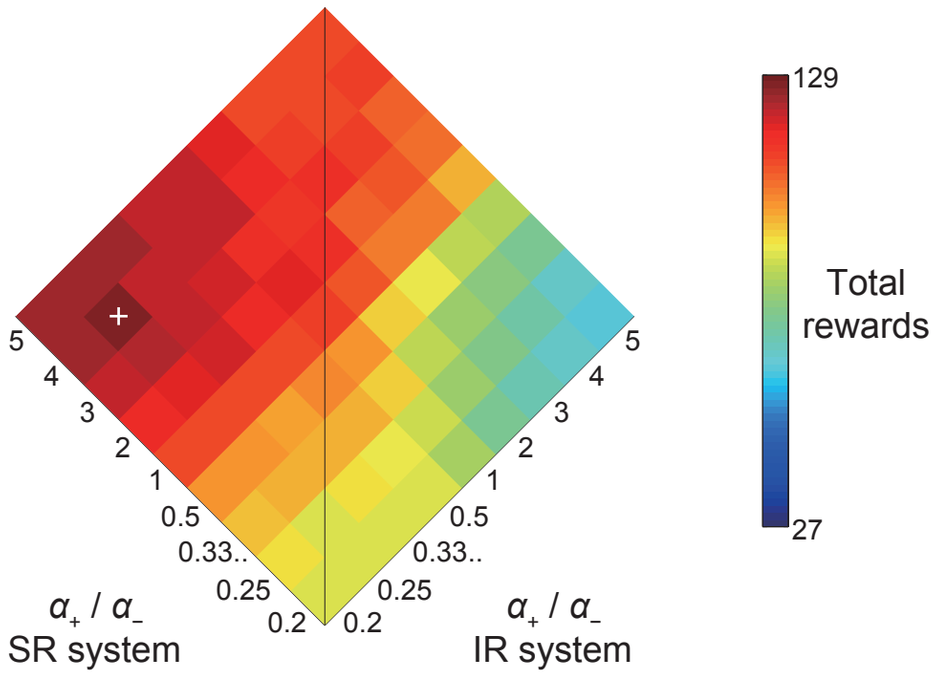
129
Total
rewards
27
5
4
3
2
1
0.5
0.33..
0.25
0.2
α_+ / α_-
SR system
α_+ / α_-
IR system

Figure 3

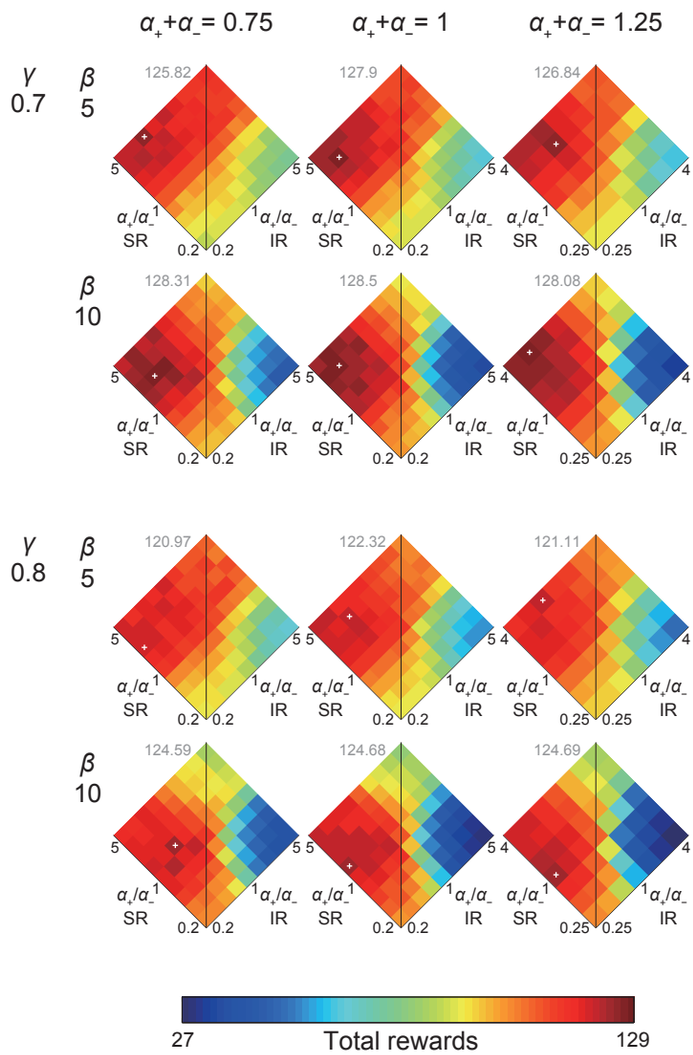
α+α- = 0.75
α+α- = 1
α+α- = 1.25
γ 0.7
β 5
β 10
γ 0.8
β 5
β 10
125.82
127.9
126.84
128.31
128.5
128.08
120.97
122.32
121.11
124.59
124.68
124.69
SR
IR
27
Total rewards
129

Figure 4

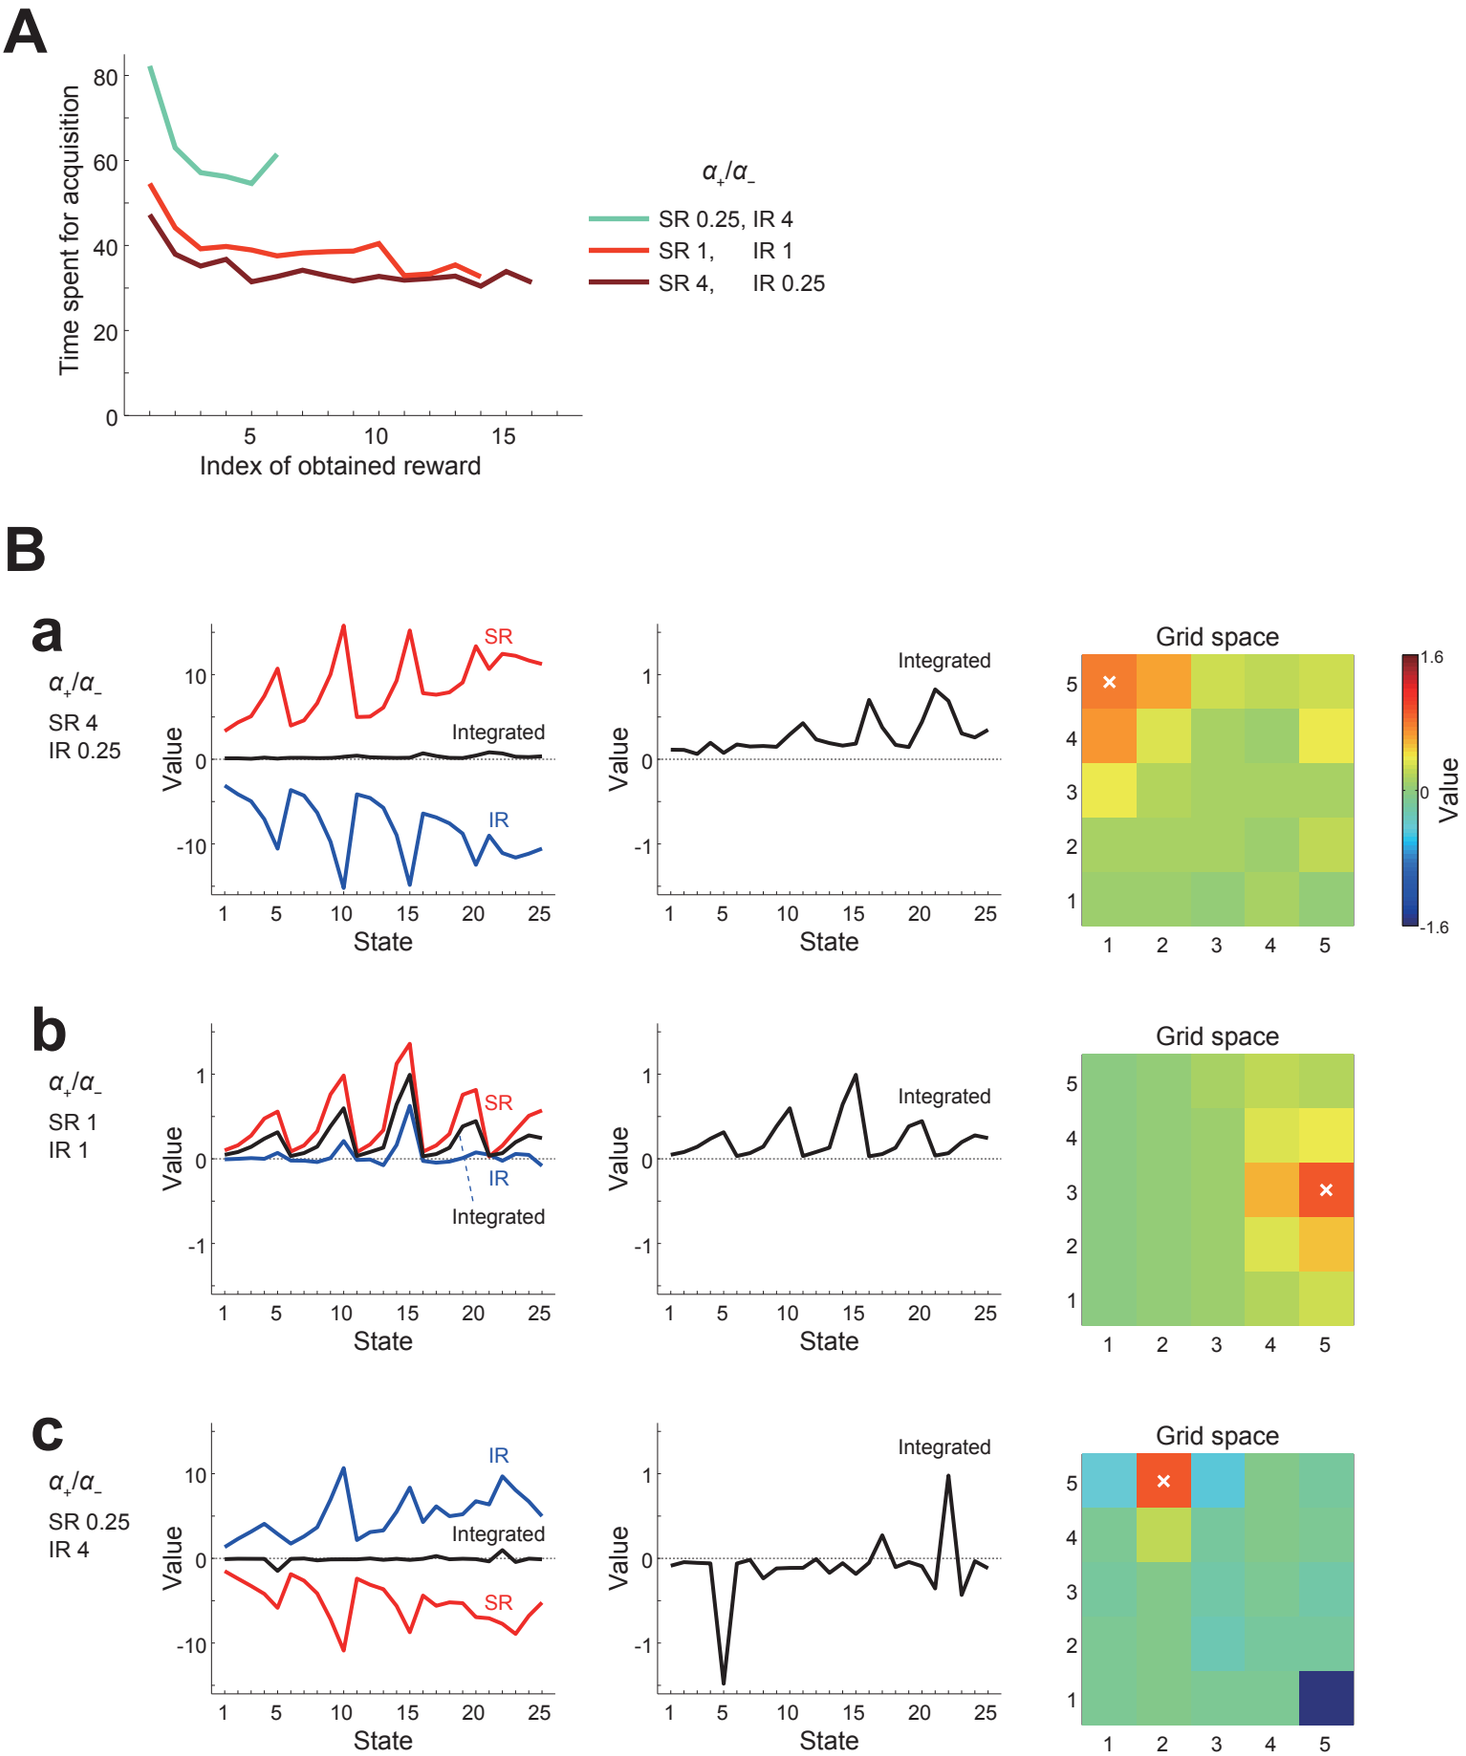
A
Time spent for acquisition
Index of obtained reward
α_+/α_-
SR 0.25, IR 4
SR 1, IR 1
SR 4, IR 0.25
B
a
α_+/α_-
SR 4
IR 0.25
SR
Integrated
IR
Value
State
Grid space
b
α_+/α_-
SR 1
IR 1
c
α_+/α_-
SR 0.25
IR 4

Figure 5

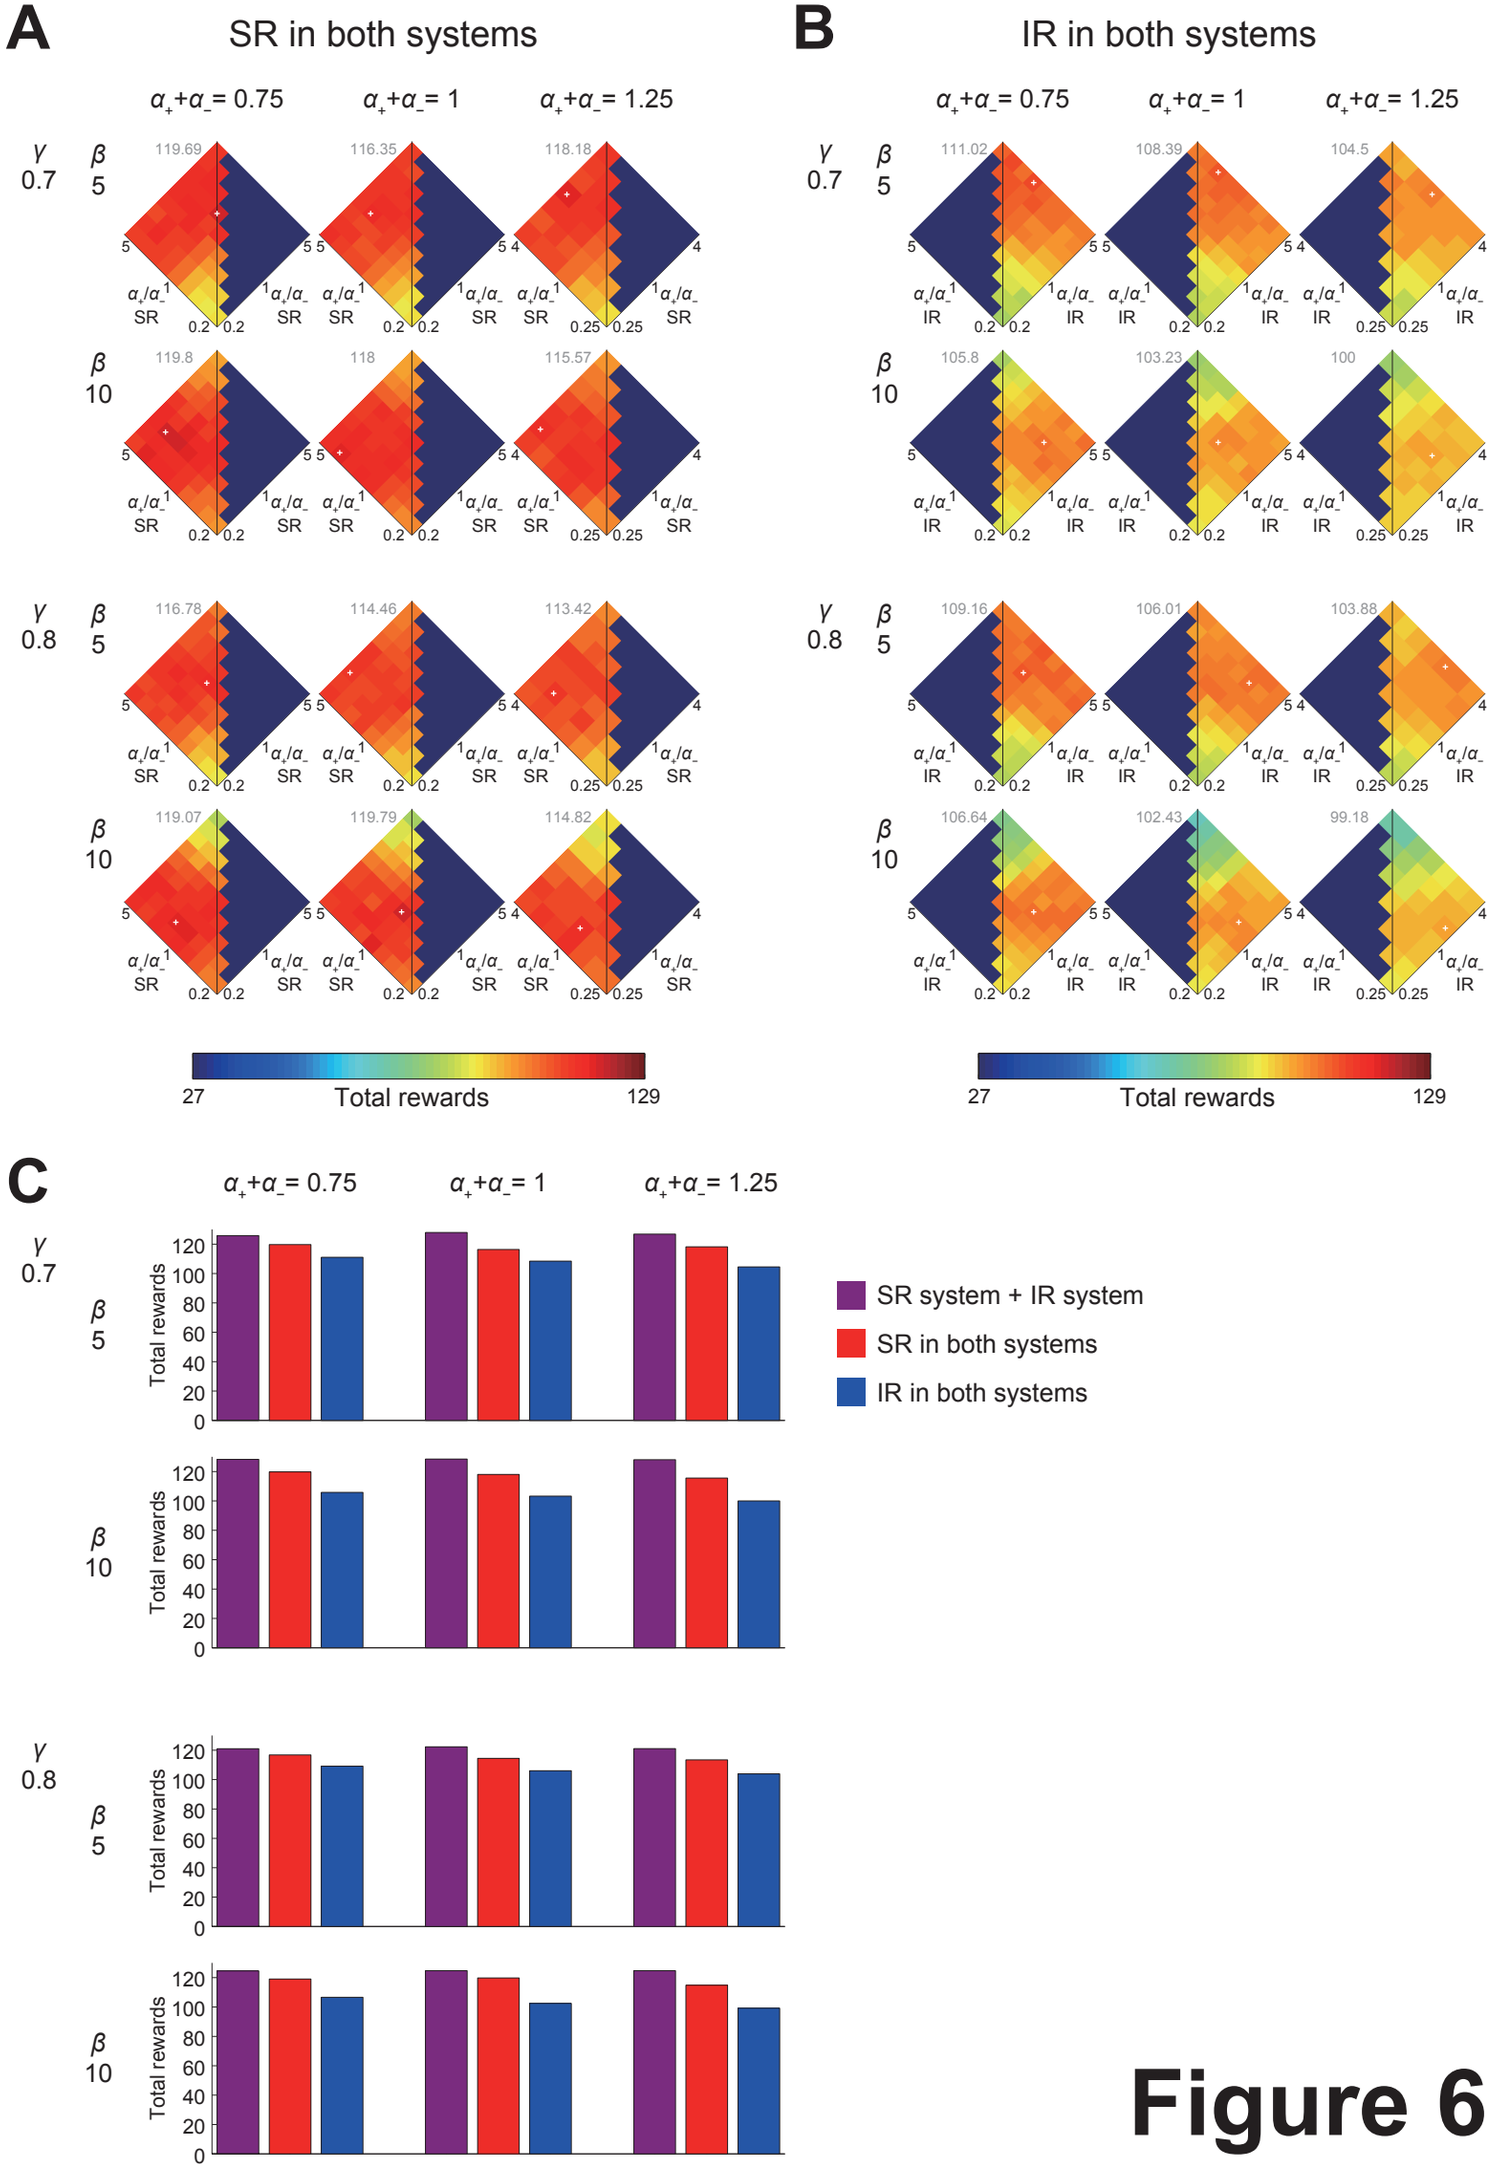
A
SR in both systems
B
IR in both systems
$\alpha_+ + \alpha_- = 0.75$
$\alpha_+ + \alpha_- = 1$
$\alpha_+ + \alpha_- = 1.25$
γ 0.7
γ 0.8
β 5
β 10
119.69
116.35
118.18
119.8
118
115.57
116.78
114.46
113.42
119.07
119.79
114.82
111.02
108.39
104.5
105.8
103.23
100
109.16
106.01
103.88
106.64
102.43
99.18
α_+/α_- SR
α_+/α_- IR
27
Total rewards
129
C
Total rewards
SR system + IR system
SR in both systems
IR in both systems

Figure 6

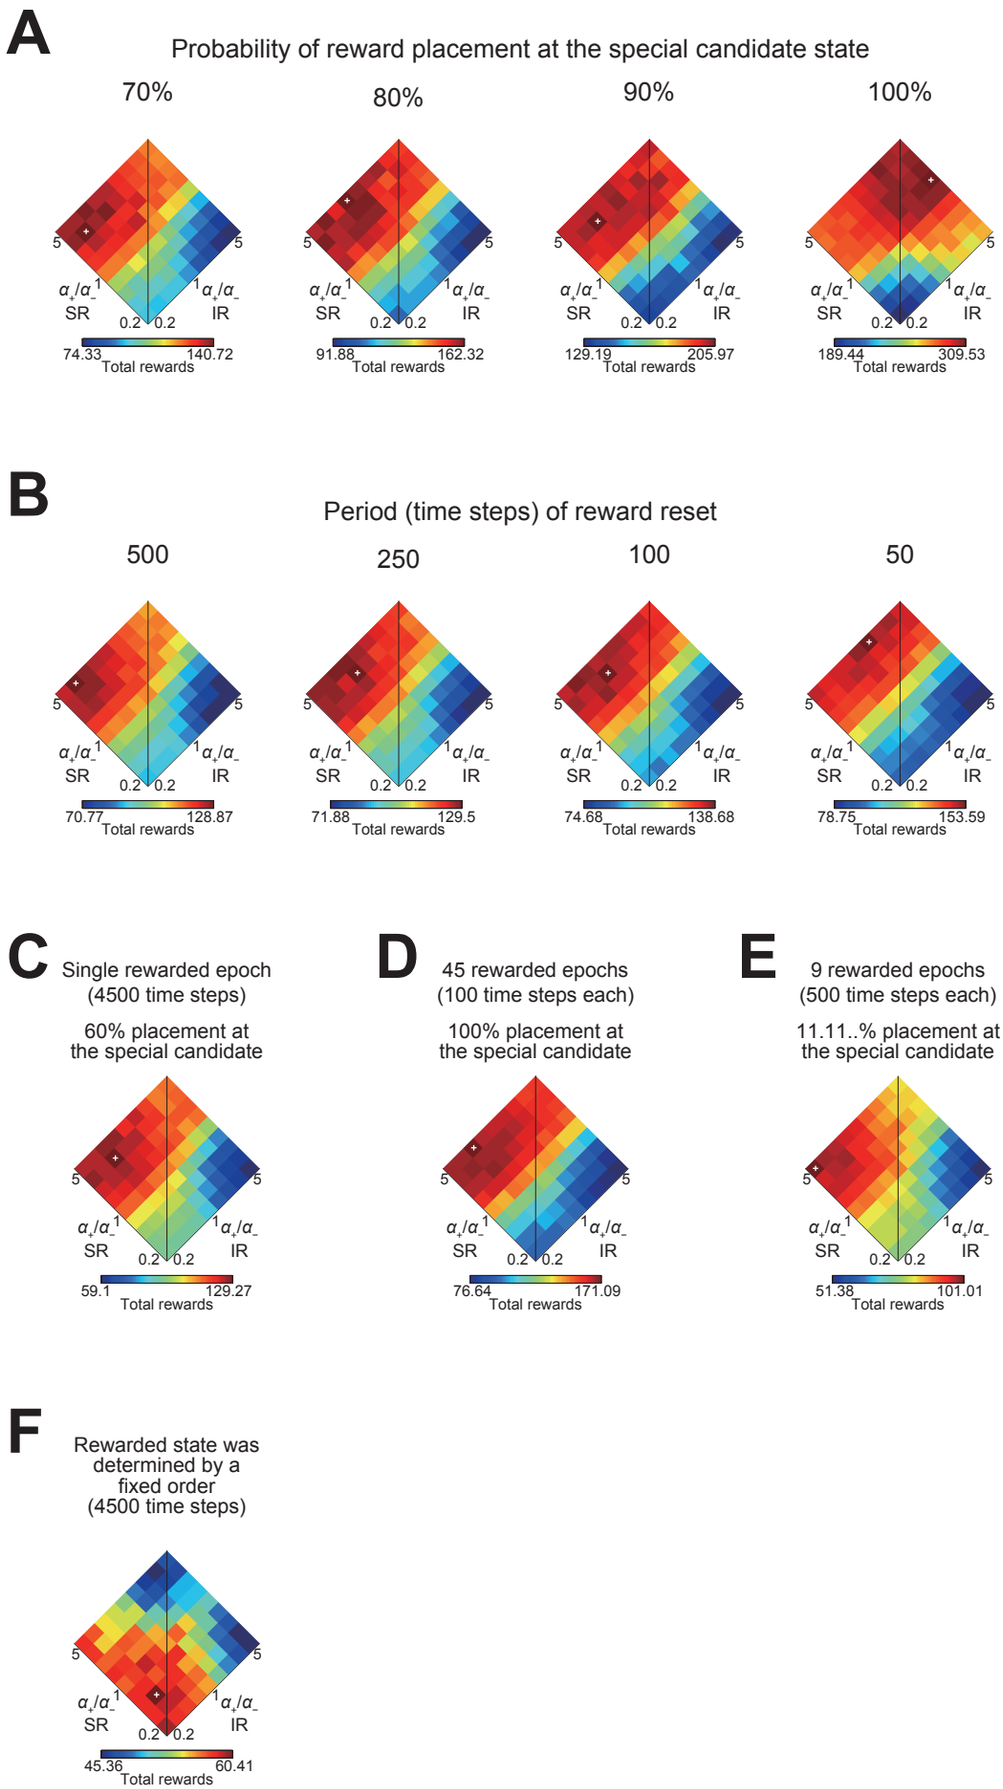
A
Probability of reward placement at the special candidate state
70%
80%
90%
100%
74.33
140.72
91.88
162.32
129.19
205.97
189.44
309.53
Total rewards
B
Period (time steps) of reward reset
500
250
100
50
70.77
128.87
71.88
129.5
74.68
138.68
78.75
153.59
Total rewards
C
Single rewarded epoch
(4500 time steps)
60% placement at
the special candidate
59.1
129.27
Total rewards
D
45 rewarded epochs
(100 time steps each)
100% placement at
the special candidate
76.64
171.09
Total rewards
E
9 rewarded epochs
(500 time steps each)
11.11..% placement at
the special candidate
51.38
101.01
Total rewards
F
Rewarded state was
determined by a
fixed order
(4500 time steps)
45.36
60.41
Total rewards
SR
IR
0.2
5

Figure 7

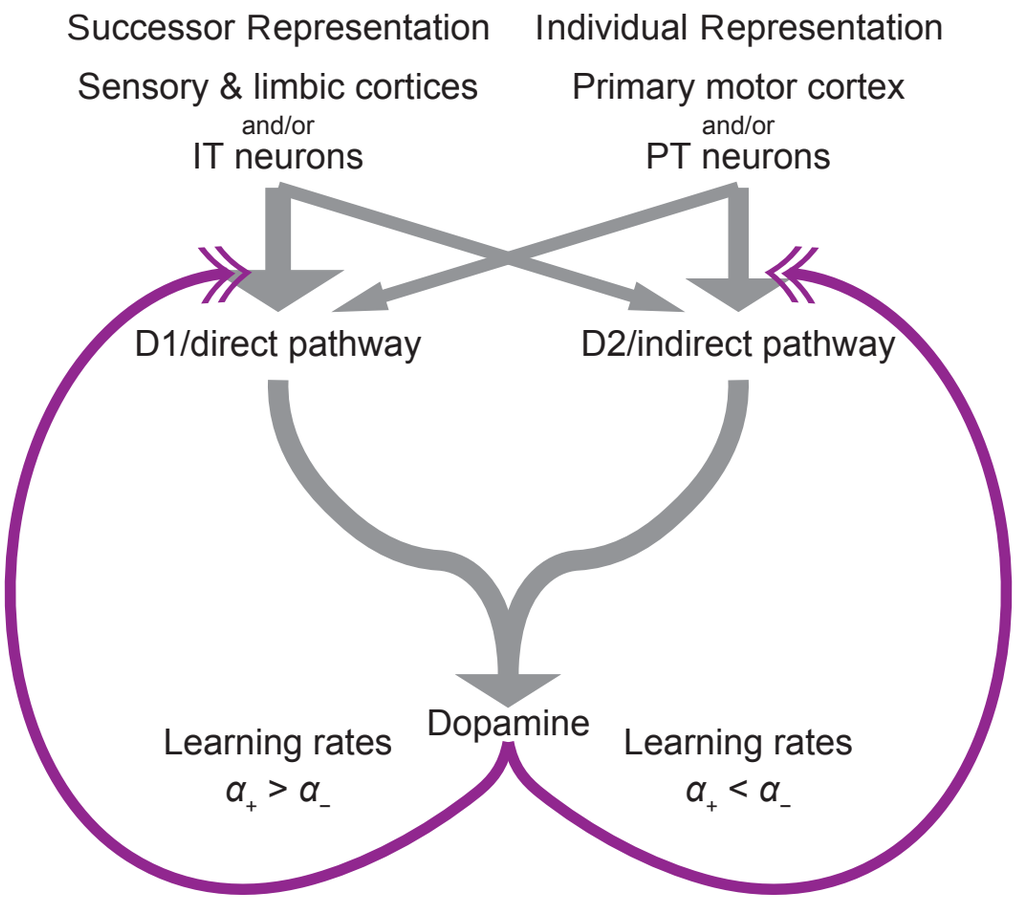
Successor Representation
Sensory & limbic cortices
and/or
IT neurons
Individual Representation
Primary motor cortex
and/or
PT neurons
D1/direct pathway
D2/indirect pathway
Dopamine
Learning rates
$\alpha_+ > \alpha_-$
Learning rates
$\alpha_+ < \alpha_-$

Figure 8