

Synergistic information in the frontal cortex-striatal pathway

Ibrahim Alsolami¹, Takashi Handa^{2,3}, and Tomoki Fukai^{1,3,*}

¹Neural Coding and Brain Computing Unit, Okinawa Institute of Science and Technology (OIST) Graduate University, Okinawa 904-0495, Japan

²Department of Neurobiology, Graduate School of Biomedical and Health Sciences (Medicine), Hiroshima University, Hiroshima 739-8527, Japan

³RIKEN Center for Brain Science, Saitama 351-0198, Japan

*Correspondence to: tomoki.fukai@oist.jp

ABSTRACT

Across the cortico-basal ganglia circuit, the medial frontal cortex (MFC) communicates with the dorsal striatum (DS) during learning and planning. How these two brain regions communicate with each other is, however, not fully understood. Here we report the presence of synergistic information during information transfer across the frontal cortex-striatal pathway. Synergistic information emerges from the positive interaction of DS and MFC neurons and provides the DS with additional cortical information. This information is held latent in neuronal signals. To reveal it, we simultaneously record neuronal activities from the MFC and DS of rats trained on an outcome-based decision-making task and determined whether past neuronal activities of the DS positively influence communication rates. We detect a neuronal synergy that enables the MFC to boost its communication rate to the DS. Our results suggest that past neuronal activities of the DS are not redundant but play a key communication role in the MFC-DS network.

Introduction

The medial frontal cortex (MFC) and dorsal striatum (DS) work in tandem to bring about positive outcomes. Across the cortico-basal ganglia circuit, the MFC projects monosynaptically to the DS^{1,2}, providing it with vital state information about the external world, such as a reward has been acquired or at times a lack thereof. In this neuronal circuit, deliberation and choice selection takes place during goal-directed behaviors³. The MFC communicates error signals to the DS and is likely to adjust action by monitoring outcomes of prior decisions⁴⁻⁶. Based on the information the DS receives from the MFC, it associates a specific selected action with a resultant outcome⁷⁻¹². This is especially evident in laboratory studies, which demonstrate that DS-lesions in rats notably compromise their performance in action selection tasks¹³.

The activity of *in vivo* cortical neurons is highly irregular, and neuronal noise limits the amount of information that can be conveyed from the MFC to the DS. Yet the DS can still make adequate behavioral decisions with such irregular and noisy cortical inputs. This unique ability of the DS likely points to a decoding mechanism operating. In this study, we explore whether such a mechanism operates in the DS to decode MFC signals. Precisely, we determine whether past neuronal activities of the DS enable it to extract more information from MFC neuronal signals than otherwise without them.

Fundamentally, the more information the DS has from the MFC, the higher the chance an adequate decision is made. A key question therefore is: how much information is the DS extracting from MFC signals? To answer this, we simultaneously recorded neuronal signals from both the MFC and DS of rats voluntarily performing a left/right licking choice task. In this task, rats had to consistently switch spouts after a reward was withheld¹⁴.

An unexpected finding in this study is that past neuronal activities of the DS, in the millisecond regime, positively influence communication rates between the MFC and DS. This finding is especially intriguing because it hints at the existence of a consecutive decoding scheme at the DS, which uses past DS signals to better decode MFC signals. Crucially, the observation that past neuronal activities of the striatum contain information needed for subsequent actions is expected^{14,15}. What is surprising here, however, is that this information is synergistic—more than merely additive (equation 5). This finding is consistent with the hypothesis that the DS is not merely decoding instantaneous MFC signals but is also actively predicting the future state of the MFC.

Results

Twelve head-restrained rats were trained to lick either a left or right spout to acquire a reward allocated to one of the spouts (Fig. 1a). The location of the rewarding spout was fixed during a trial block and was reversed after a rat accumulated a total of 10 rewards (Fig. 1b). During this behavioral task, we simultaneously recorded neuronal activity from both the MFC and DS on the left hemisphere of rats with a pair of silicon probes. A retrograde tracer confirmed the projection from the recording site at the MFC to its corresponding recording site at the DS (Fig. 2). In-depth details of our experiment are available in our recent report¹⁴.

Synergistic Information and its Implications in Neural Communication

We quantify the amount of synergy forming between the MFC and DS from an information-theoretic standpoint. We briefly explain the information-theoretic measures employed in this study. In Fig. 3, X is a random variable (RV) that counts the number of spikes generated by all MFC neurons within a time-window of duration τ_1 . Likewise, Y and Z are RVs that count the number of spikes generated by all DS neurons within a non-overlapping sliding time-window of duration τ_1 and τ_2 , respectively. Herein, we evaluate the amount of synergistic information at the neural population level but not at that of single neurons, where large noise may arise due to sparse neuronal firing.

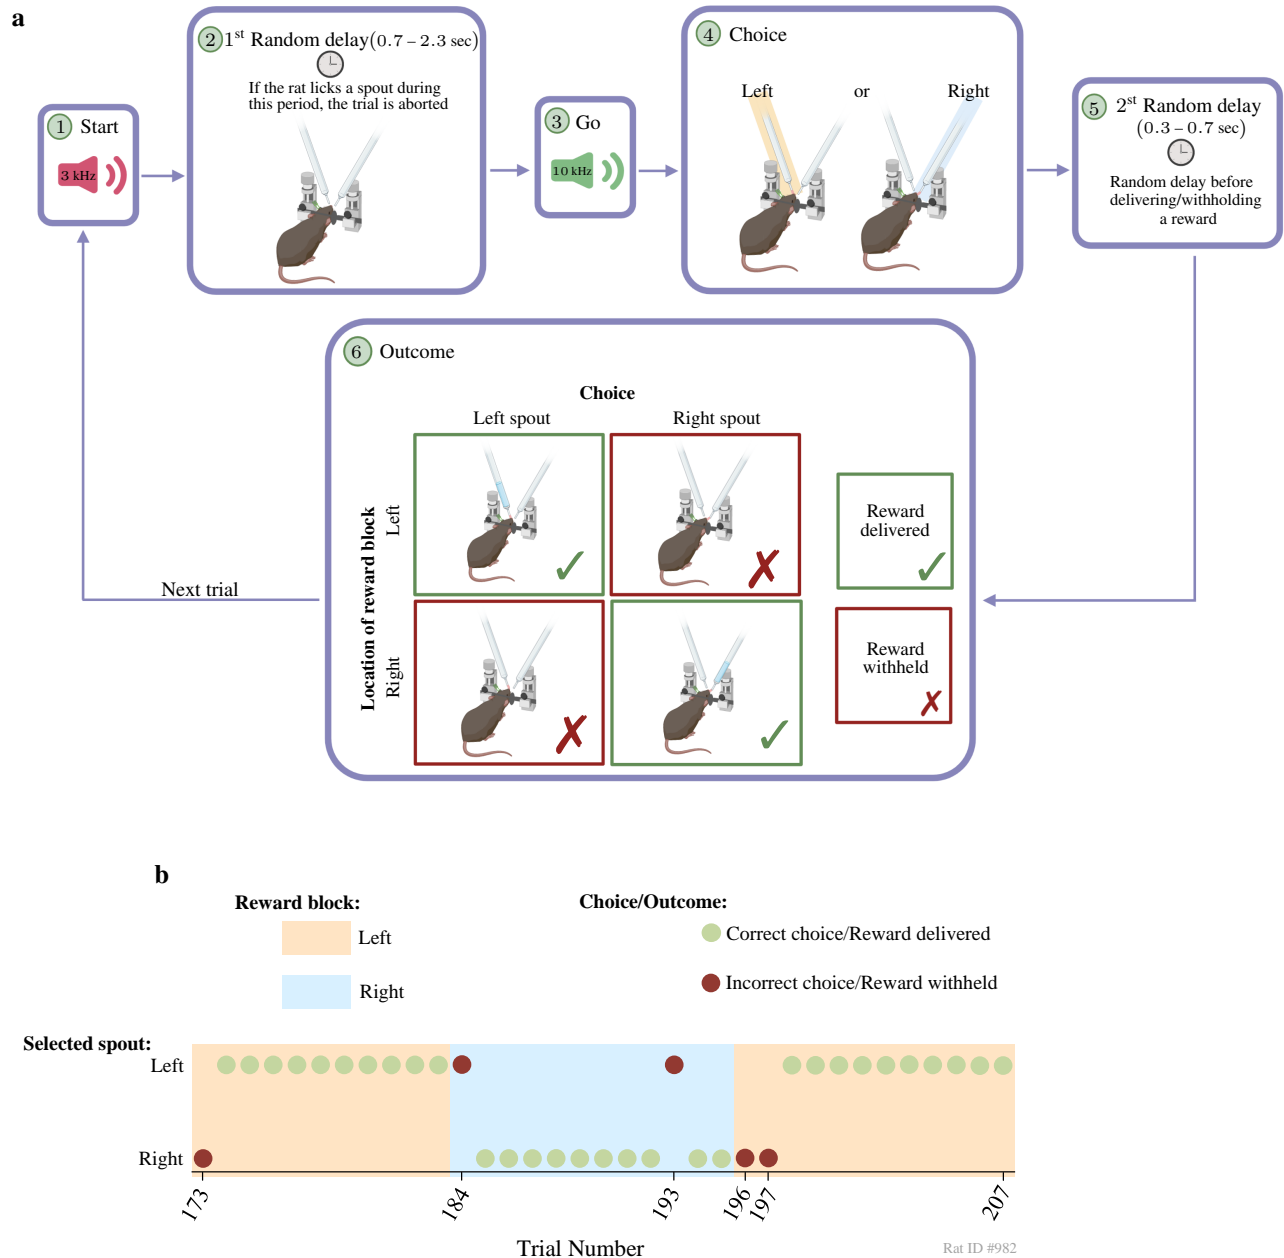


Figure 1. Experiment. **a** Schematic illustration of a single test trial: (1) A trial starts with a 3 kHz (1 sec) tone. (2) A random delay (0.7–2.3 sec) is then introduced. If a rat licks a spout during this delay period, the trial is aborted. (3) Next, a 10 kHz (0.1 sec) “Go” signal is presented. (4) After this signal, a rat selects either a left or right spout. (5) This is followed by a second random delay (0.3–0.7 sec). (6) Lastly, a reward is delivered if the selected spout location matches the reward block location; otherwise, a reward is withheld. After this step, a new trial begins. **b** An example from our behavioral experiment. First, the rat selects a spout (left/right). Then, a reward is delivered if the rat selects a spout location that matches the location (left/right) of a reward block. Otherwise, the reward is withheld. The location of the reward block is reversed after the rat accumulates a total of 10 rewards. Throughout the experiment, we provided no sensory feedback to rats regarding the location of rewards. In the above example, all the decisions made by the rat are correct except for trials #173, 184, 193, 196, and 197 (because the location of the selected spout does not match the location of the reward block). The incorrect choices made in trials #173, 184, and 196 are highly expected as rats tend to select the same rewarding spouts until it does not provide a reward. After which, rats tend to switch to an opposite spout, such as the decisions illustrated in trails #174, 185, 194, and 198 (that is to say, rats exhibit a win-stay/lose-shift strategy). In trial #197, however, the rat made a consecutive incorrect choice despite not gaining a reward in the previous trial #196; this case was not uncommon in our experiment. Additionally, in trial #193, the rat switched to an incorrect spout despite gaining a reward in the previous trial #192; while this case is uncommon, we observed it in our experiment.




Figure 2. Fluorescent images of brain sections. Bottom left: Fluoro-gold labelled MFC neurons and injection site of tracer (Fluoro-Gold) at the DS. Top right: Expanded view of labelled MFC neurons in deep layers.

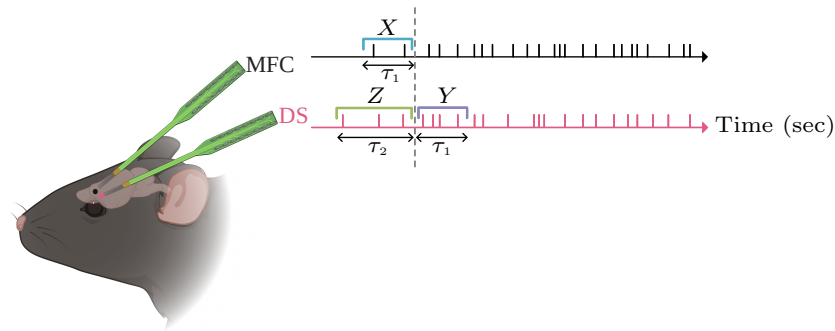


Figure 3. An illustration of the electrophysiological recording and random variables (RVs) involved. The MFC spike train displays the collective firing of all MFC neurons, rather than the firing of an individual MFC neuron, and likewise for the DS spike train. RV X counts the number of spikes generated by all MFC neurons within a time-window of duration τ_1 . Similarly, RVs Y and Z count the number of spikes generated by all DS neurons within a time-window of duration τ_1 and τ_2 , respectively.

The mutual information quantifies the amount of information shared between two RVs and is defined as^{16,17}

$$I(X;Y) = H(X) - H(X|Y) \text{ bits}, \quad (1)$$

where $H(X)$ is the entropy of X and $H(X|Y)$ is conditional entropy of X given Y . The mutual information can be equivalently interpreted as the amount of uncertainty remaining about X after observing Y . The mutual information, however, satisfies $I(X;Y) = I(Y;X)$, implying that it does not describe the directional flow of information between two processes.

Unlike mutual information, the transfer entropy captures the directional flow of information between X (e.g., neuronal activity in the MFC) and Y (e.g., neuronal activity in the DS)^{18–21}, and is defined as

$$I(X;Y|Z) = H(X|Z) - H(X|Y,Z) \text{ bits}, \quad (2)$$

with Z being the history of Y (Fig. 3). Moreover, $H(X|Z)$ is the conditional entropy of X given Z , and $H(X|Y,Z)$ is the conditional entropy of X given both Y and Z .

In this study, we test whether past neuronal activities, Z , give the DS the ability to extract more information from MFC signals than otherwise without them. For this purpose, we need to measure the interaction influence of Z . We can quantify this influence by computing the interaction information, S , which is defined as the extra transfer entropy remaining after subtracting the mutual information^{22–24}:

$$S = I(X;Y|Z) - I(X;Y) \text{ bits}. \quad (3)$$

If $S > 0$, then Z positively influences communication rates and accordingly yields synergistic information (equation 5). Otherwise, if $S < 0$, then Z gives redundant information. With some algebraic manipulation, we can express equation 3 as

$$S = I(X;Y,Z) - (I(X;Y) + I(X;Z)) \text{ bits}. \quad (4)$$

When $S > 0$, we have

$$I(X;Y,Z) > I(X;Y) + I(X;Z) \text{ .} \quad (5)$$

In words, equation 5 says that the whole, $I(X;Y,Z)$, is greater than the sum of the parts, $I(X;Y) + I(X;Z)$, or equivalently a synergy is present—i.e., when Y and Z work together to decode X , more information about X is obtained than if they were to work individually. As such, the interaction information reveals whether knowing Z enables us to extract additional information from Y about X .

Positive interaction information can convey an important message about the neuronal communication between the MFC and DS. When Z represents the history of Y in the immediate past, S is likely to be positive if the neuronal activities of the DS (Y and Z) can, at least to some extent, predict the neuronal activity of the MFC (X). This hypothesis is worthwhile to examine: if the hypothesis is true, the DS is not only processing the current state of the MFC during decision-making but is also predicting future cortical states.

Presence of Synergistic Information

Figure 4 shows rates of synergistic information (bits/sec) obtained from the neuronal activity data of 12 rats in four different trial conditions (reward versus non-reward, and left choice versus right choice). Here, we report rates of synergistic information (\bar{S}/τ_1) as opposed to synergistic information alone (\bar{S}) so as to put quantities in a time context. For the reason mentioned later in this section, we examined synergistic information rates for all possible combinations of reward locations (left or right spout) and outcomes (reward given or withheld). The magnitudes of the rates significantly varied from rat to rat due to certain reasons such as differences in sample sizes and average firing rates of neurons. In the majority of cases (at the $p = 0.05$ level), a synergistic information rate is present, that is, $\frac{\bar{S}}{\tau_1} > 0$. However, in Fig. 4c for rodent ID #1012 and at $\tau_2 = 10$ ms a faint synergy was observed but not significant at the $p = 0.05$ level and therefore was discarded (and similarly, in Fig. 4d, for rat #982 and at $\tau_2 = 10$ and 20 ms; as well as in Fig. 4d for rat #1012 and at $\tau_2 = 20$ ms). Additionally, Fig. 4 reveals that, by and large, there is a gradual increase in rates of synergistic information as more history, τ_2 , is included.

We examined synergistic information rates during a behavioral task to rule out the possibility that the presence of synergistic information might be reward reliant or hemisphere dependent. The reason being is that in our experiment, silicon probes were inserted into rats' left hemisphere. As such, in the behavioral task, we switched the location of rewards intermittently from left to right and vice versa, giving a total of 4 possible combinations: left/right and reward/non-reward. In our analysis, we set τ_1 to 5 ms as this value is within the range of time it takes an MFC signal to reach the DS¹. Additionally, we vary the value of τ_2 across a range of values (5 – 20 ms) because a small value of τ_2 may risk under-estimating the transfer entropy, while a large value of τ_2 may, on the other hand, risk over-estimating the transfer entropy as a result of under-sampling the underlying probability distribution function²⁰. We emphasize that in nearly all cases, synergistic information is present.

Finally, we test whether synergistic information in the DS is related to the sequentiality of population activities observed in DS neurons¹⁵. Neuronal populations in the MFC and DS often exhibit similar sequential firing patterns during behavioral tasks. Recently, a notable difference between MFC and DS activity patterns was observed in rodents performing a two-interval timing task¹⁵, and an outcome-based alternative choice task (in which the present data were recorded)¹⁴. In both studies, neuronal dynamics of the DS showed a higher degree of sequentiality than those of the MFC. It is hypothesised that the higher the sequentiality of a neuronal signal, the easier it is for downstream brain areas to monitor and read out time from such signal. In the DS, monitoring time could potentially imply that it is tracking cortical inputs. However, no significant correlation was found in the present study between synergistic information rates and sequentiality across rats (Supplementary Fig. 1).

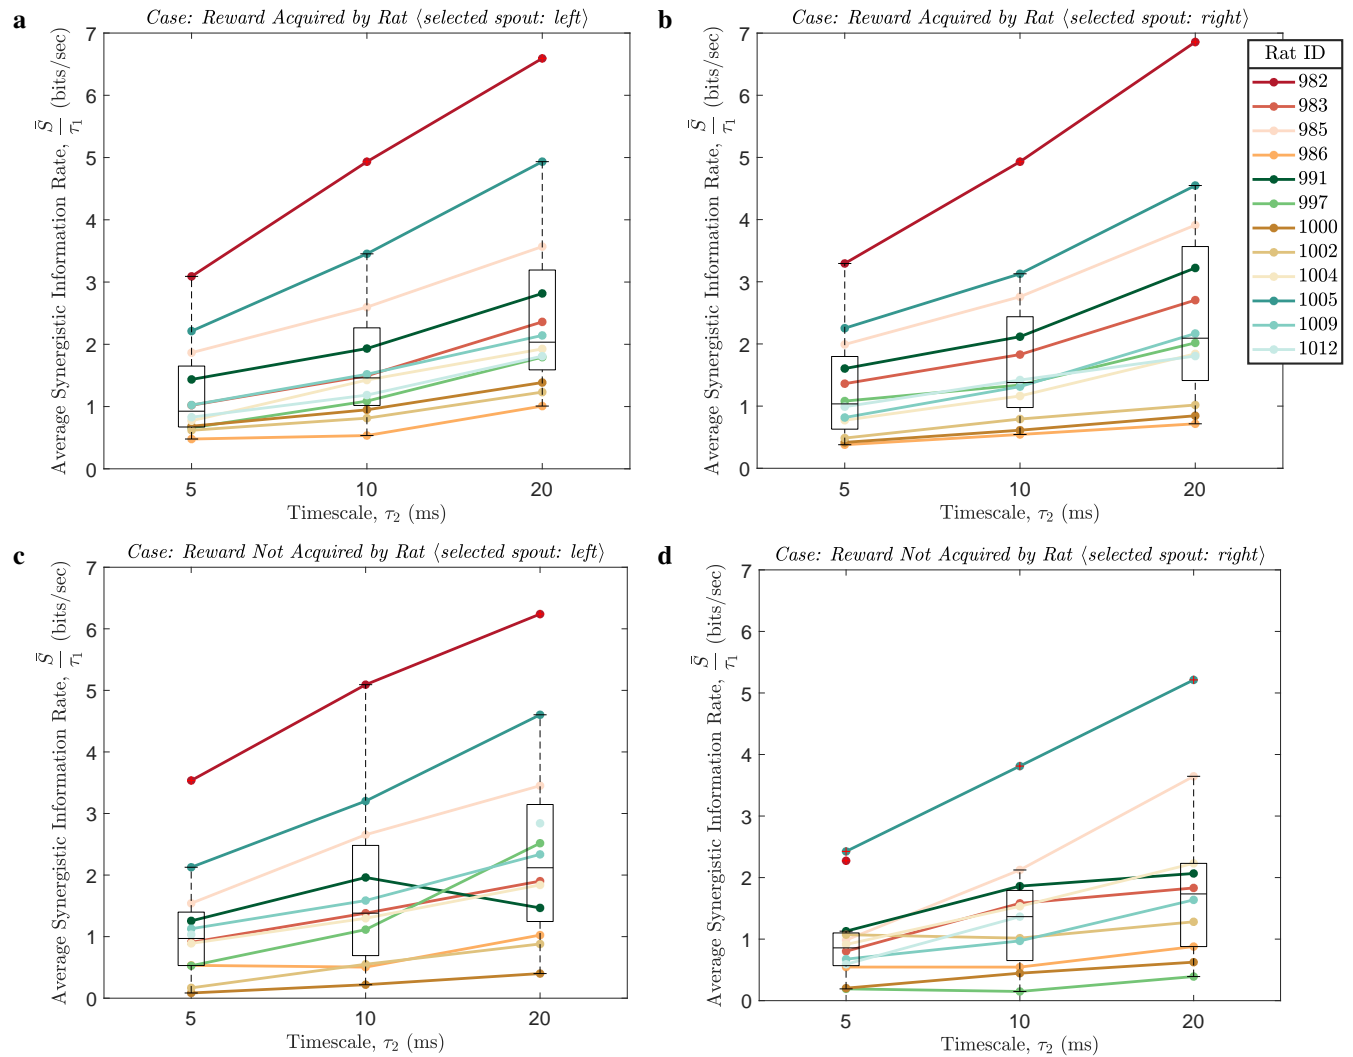


Figure 4. Synergistic information rates. **a** Reward acquired by rat, and the selected spout is left. **b** Reward acquired by rat, and the selected spout is right. **c** Reward not acquired by rat, and the selected spout is left. **d** Reward not acquired by rat, and the selected spout is right. In Fig. 4c, for rat #1012 at $\tau_2 = 10$ ms a faint synergy was observed but not significant at the $p = 0.05$ level and therefore was discarded (and similarly, in Fig. 4d for rat #982 at $\tau_2 = 10$ and 20 ms; as well as in Fig. 4d for rat #1012 at $\tau_2 = 20$ ms). Here we set τ_1 to 5 ms as this value is within the range of time it takes an MFC signal to reach the DS¹. Additionally, we vary the value of τ_2 from 5 to 20 ms to observe its influence on information rates. Boxes within plots represent the 25th–75th percentiles and middle lines of boxes mark median values. Additionally, outlier points are beyond the ends of boxplots' whiskers and marked with a '+' symbol.

Discussion

A synergistic information rate is present as past neuronal activities of the DS (Z) enables it to extract more information ($I(X;Y|Z)$) from MFC signals than otherwise without them ($I(X;Y)$). What is interesting about the data is that the DS can decode information contained in MFC signals in a synergistic ($I(X;Y,Z)$) rather than an additive manner ($I(X;Y) + I(X;Z)$). Our results suggest that the history of DS activities plays a crucial role in the MFC-DS communication network because it enhances the communication rate between the two brain regions. Whether past striatal activities contribute to the cortico-striatal information transfer has not been previously explored.

Searching for synergistic information in a neural code has been undertaken and found in a variety of settings. In blowflies, synergistic information was detected in H1-neurons. It was observed that two action potential spikes together, across time, carry more than twice the amount of information of individual spikes²³. In salamanders and guinea pigs, a similar synergy was detected in a combinatorial code of neurons, whereby spiking and silent neurons jointly produce synergistic information in a neural pattern²⁵. Likewise, synergistic information in patterns of neuronal activities at the population level was further observed in macaque monkeys²⁶. The previously described studies examined synergistic information in neuronal representations within a single brain region. What sets our work apart from prior research work is that we examine the synergy forming across brain regions—the MFC and DS, the critical brain regions implicated in decision-making. Our results indicate that synergistic information is also present in neural codes for cross-area communication.

The interaction between the MFC and DS has been gaining attention. Results from earlier studies revealed that ramping activities of the MFC and DS are likely to represent temporal signals that encode the passage of time²⁷. Additionally, it was demonstrated that optogenetically stimulating MFC-to-DS axonal projections increase time-dependent ramping in the DS²⁸. In these studies, the linear relationship between the MFC and DS was studied. Our results are complementary to these studies in that we examined a nonlinear relationship between the MFC and DS by means of information-theoretic tools. As such, a novel interaction is observed, namely, the emergence of synergistic information once past neuronal activities of the DS are considered.

We propose that synergistic information, at least partially, emerges from the computational effort of the DS to predict neuronal activity in the MFC in the immediate future—within the milliseconds range. We consider that this predictive role of the DS is unrelated to reward prediction. A body of evidence shows that the striatum represents error signals of reward predictions^{29,30} that modify the efficacy of neuronal signal transmission in the cortico-striatal pathway^{31,32}. The timescale of this prediction is behaviorally relevant in the range of seconds, which is much shorter than the millisecond range timescale of synergistic information at the DS. Additionally, no particular differences between rewarded and unrewarded trials in terms of rates of synergistic information were observed in our experiment.

Synergistic information gained by past DS activities is uncorrelated with the degree of sequentiality of the population activity in the DS, suggesting that monitoring time and predicting cortical states are not conclusively related to each other in the present study. This may simply reflect the fact that unlike the interval timing task¹⁵, the present outcome-based choice task does not require an accurate estimation of time. However, we also noticed that the sequentiality index used in this and previous studies is not robust against noise. The computational relationship between synergistic information and high sequentiality in the DS and the behavioral implications of these observations require further clarifications.

From an energy perspective, synergistic information improves metabolic efficiency. In recent years, there has been a particular interest in studying the energy efficiency of the mammalian brain. One line of research indicates that the mammalian brain is structured in such a way as to maximize metabolic efficiency^{33–36}—measured as bits transmitted per energy expended to generate an information-carrying action potential. In other words, the brain is tuned to communicate economically in terms of energy expended to transmit information. Our results are in line with the metabolic efficiency viewpoint because past neuronal activities enable the DS to extract more information contained in MFC signals than without them. This, as a result, increases the metabolic efficiency of the cortico-striatal circuit: more bits extracted from a neuronal signal gives a higher metabolic efficiency. Additionally, it is useful to mention that the efficient use of metabolic energy is a widely accepted assumption behind the development of the so-called predictive coding hypothesis for hierarchical cortical computing^{37–39}, although counter-arguments have also been given⁴⁰.

A suggested hypothesis for the function of the basal ganglia proposes that the striatum selectively regulates (gates) cortical information. When the DS regulates cortical information, task-relevant information has to be extracted from MFC signals and subsequently transmitted to downstream brain regions. Such processing of cortical signals will filter out information unnecessary for controlling voluntary movements^{41–44}. Additionally, filtering at the DS is more likely to decrease rather than increase the amount of information extracted from the MFC signals. Our results, however, suggest that more information flows from the MFC to DS when past neuronal activities of the DS are considered than without them (i.e., $I(X;Y|Z) > I(X;Y)$ or equivalently $S > 0$, equation 3). Although the gating hypothesis and the existence of synergistic information at the DS are not mutually exclusive, they do not harmonize with each other.

A shortcoming of our study is that we only observed a small area of the MFC and DS using two probes. We used a retrograde tracer to confirm the projection from the recording site at the MFC to its corresponding recording site at the DS. However, a large-scale experiment covering a wide area of both the MFC and DS is necessary to solidify our results. Nonetheless, all the rats examined in this study consistently supported the significant levels of synergistic information at the DS. We, therefore, believe that our results are not mere noise arising from the limited sampling of MFC and DS neuronal signals but indicate a novel predictive function operating in the frontal cortex-striatal pathway on short timescales.

Methods

Animal preparation

All procedures followed in this study were carried out in accordance with the Animal Experiment Plan, which was reviewed and approved by the Animal Experiment Committee of RIKEN. A total of 20 Long-Evans rats (6 weeks, male, 200–220 g, Japan SLC Inc.) were used in our experiment. In-depth details of our methods are available in our recent report¹⁴. In the present study, we analyse the data obtained from our recent work¹⁴.

Behavioral task

We used a customized multiple-rat training system^{14,45} to train several rats in parallel to learn the spout-selection task and gave them saccharin-water (0.1%, 15 μ l) as a reward. We analysed the neuronal activity of 12 rats out of the 20 we trained because they provided a satisfactory number of units from both the MFC and DS.

Figure 1a provides a schematic illustration of a single test trial. A trial starts with a 3 kHz (1 sec) tone. Then, a random delay (0.7–2.3 sec) is introduced; if a rat licks a spout during this period, the trial is aborted. The random delay is followed by a 10 kHz (0.1 sec) “Go” signal. After this signal, a rat selects either a left or right spout. Then, a second random delay (0.3–0.7 sec) is introduced before a fluid reward is either delivered or withheld; if a rat selects a correct spout location, a reward is delivered for a period of 4 sec. On the other hand, if a rat selects an incorrect spout location, a reward is withheld; the duration of this non-reward event is 5 sec. After a reward/non-reward phase is complete, a new trial begins.

Figure 1b provides an example from the behavioral experiment. First, a rat selects either a left or right spout. Then, a reward is delivered if the rat selects a correct choice. A choice is declared correct if the selected spout location matches the location of a reward block. Otherwise, the choice is declared incorrect, and a reward is withheld. The location of the rewarding spout was fixed during a trial block and was reversed after a rat accumulated a total of 10 rewards. During the experiment, we provide no sensory feedback to the rats regarding the location of the reward.

Electrophysiological recordings from the MFC and DS

Figure 3 shows an illustration of the electrophysiological recording. We inserted the pair of silicon probes (32-channels, 4 shanks separated 0.4 mm apart) that have tetrode-like electrodes (A4 \times 2-tet-7/5 mm 500-400-312, NeuroNexus Technologies) as follows:

- The first silicon probe was inserted vertically (depth from the pia mater: 1.2 mm) into the MFC (+2.4–4.2 mm from the Bregma and 1.0–1.4 mm to the midline).
- The second silicon probe was inserted at an angle of 6° posteriorly into the DS through a cranial window (+0.6–1.0 mm from the Bregma and 1.0–1.4 mm to the midline).

Retrograde tracing

To confirm that the recording site at the MFC projects to that of the DS, we injected a retrograde tracer (Fluoro-Gold) into the DS 3 days before a first electrophysiological recording session and checked the position of labelled neurons *post hoc*. Fig. 2 shows fluorescent images obtained after a recording experiment. We observed that corticostriatal projection neurons at the MFC are labelled with fluoro-gold after injecting the retrograde tracer at the DS. We collected images with a fluorescent microscope (Olympus AX70, Tokyo, Japan).

Measurements and calculations

Neuronal signals of our experiment are analysed from the moment a rat first licks a spout (after the onset of a “Go” signal) to the end of a trial (Fig. 1a). During this period, the probabilities of events and synergistic information are computed as follows:

- **Probabilities:** Fig. 3 provides an illustration of the random variables (RVs) and time-windows considered herein.

Let X be an RV that counts the number of spikes generated by all MFC neurons within a time-window of duration τ_1 . Likewise, let Y and Z be RVs that count the number of spikes generated by all DS neurons within a time-window of duration τ_1 and τ_2 , respectively.

Now let us describe how outcomes of RVs are obtained. In Fig. 3, all three time-windows jointly slide by increments of τ_1 to obtain an event (x, y, z) , where x , y , and z are outcomes of X , Y , and Z , respectively. Namely, x , y , and z are the number of spikes observed in time intervals $[n\tau_1, (n+1)\tau_1]$, $[(n+1)\tau_1, (n+1)\tau_1 + \tau_2]$, and $[n\tau_1, (n+1)\tau_1]$, respectively; where $n \in \{0, 1, 2, \dots\}$ (see Supplementary Fig. 2 for a detailed illustration).

Probabilities are estimated by the relative frequency of occurrence of events. For instance, $p(x, y, z)$ is the relative frequency of occurrence of the event $(X = x, Y = y, Z = z)$.

- **Synergistic information:** The amount of synergistic information is computed as follows:

$$S = I(X; Y|Z) - I(X; Y) \quad \text{bits}, \quad (6)$$

where $I(X; Y|Z)$ is the transfer entropy^{18–21} and $I(X; Y)$ is the mutual information^{16,17}, and are calculated as follows:

$$I(X; Y|Z) = \sum_{x,y,z} p(x, y, z) \log_2 \frac{p(x, y, z)p(z)}{p(x, y)p(y, z)} \quad \text{bits} \quad (7)$$

and

$$I(X; Y) = \sum_{x,y} p(x, y) \log_2 \frac{p(x, y)}{p(x)p(y)} \quad \text{bits}. \quad (8)$$

Statistical significance

This section is based in part on^{46–49}. We test the statistical significance of our results by the following procedure:

- First, for each experimental trial, we generate null (surrogate) data of the MFC by the following method: consider the illustration shown in Supplementary Fig. 3. The MFC spike train of each experimental trial is divided into two parts by a random number, t^* , drawn from a uniform distribution $U(t_{\text{start}}, t_{\text{end}})$, where t_{start} and t_{end} are the start and end time of the MFC spike train, respectively. Next, the two MFC parts are swapped to create null data. Ideally, this makes the MFC and DS spike trains independent of each other—effectively disconnecting the MFC from the DS in terms of shared information while preserving much of the statistical properties of the original MFC spike train. This randomization procedure is repeated $N = 10^4$ to form null data for each experimental trial. Additionally, the DS spike train was left unrandomized to preserve the statistical relationship between Y and Z ⁴⁹.
- Second, using the MFC null data generated in the previous step along with the intact DS spike train, we use equations 7 and 8 to compute the set $I_s(X;Y)$ and $I_s(X;Y|Z)$ (the null hypothesis versions of $I(X;Y)$ and $I(X;Y|Z)$, respectively).
- Third, the p -value of $I(X;Y)$ is estimated as the fraction of $I_s(X;Y)$ values that are greater or equal to $I(X;Y)$. Likewise, the p -value of $I(X;Y|Z)$ is estimated as the fraction of $I_s(X;Y|Z)$ values that are greater or equal to $I(X;Y|Z)$.

Calculation remarks

- After we calculate $I(X;Y)$ and $I(X;Y|Z)$, and their associated p -value for each experimental trial, we take the average of all S values for which both $I(X;Y)$ and $I(X;Y|Z)$ have a p -values ≤ 0.05 . We denote this average by \bar{S} .
- Consider Fig. 1b. If a trial before an incorrect choice is as well incorrect, then it is discarded in our calculations. For example, trial #196 in Fig. 1b was discarded while trial #197 was retained. Such events are not uncommon in our experiment, and we discard them to preserve consistency.

References

1. Wilson, C. J. Postsynaptic potentials evoked in spiny neostriatal projection neurons by stimulation of ipsilateral and contralateral neocortex. *Brain research* **367**, 201–213 (1986).
2. McGeorge, A. & Faull, R. The organization of the projection from the cerebral cortex to the striatum in the rat. *Neuroscience* **29**, 503–537 (1989).
3. Lo, C. C. & Wang, X. J. Cortico–basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat. neuroscience* **9**, 956–963 (2006).
4. Narayanan, N. S., Cavanagh, J. F., Frank, M. J. & Laubach, M. Common medial frontal mechanisms of adaptive control in humans and rodents. *Nat. neuroscience* **16**, 1888–1895 (2013).

5. Hyman, J. M., Holroyd, C. B. & Seamans, J. K. A novel neural prediction error found in anterior cingulate cortex ensembles. *Neuron* **95**, 447–456 (2017).
6. Bonini, F. *et al.* Action monitoring and medial frontal cortex: leading role of supplementary motor area. *Science* **343**, 888–891 (2014).
7. Balleine, B. W., Delgado, M. R. & Hikosaka, O. The role of the dorsal striatum in reward and decision-making. *J. Neurosci.* **27**, 8161–8165 (2007).
8. Nicola, S. M. The nucleus accumbens as part of a basal ganglia action selection circuit. *Psychopharmacology* **191**, 521–550 (2007).
9. Reynolds, J. N., Hyland, B. I. & Wickens, J. R. A cellular mechanism of reward-related learning. *Nature* **413**, 67–70 (2001).
10. Lauwereyns, J., Watanabe, K., Coe, B. & Hikosaka, O. A neural correlate of response bias in monkey caudate nucleus. *Nature* **418**, 413–417 (2002).
11. Isomura, Y. *et al.* Reward-modulated motor information in identified striatum neurons. *J. Neurosci.* **33**, 10209–10220 (2013).
12. Nonomura, S. *et al.* Monitoring and updating of action selection for goal-directed behavior through the striatal direct and indirect pathways. *Neuron* **99**, 1302–1314 (2018).
13. Skelin, I. *et al.* Lesions of dorsal striatum eliminate lose-switch responding but not mixed-response strategies in rats. *Eur. J. Neurosci.* **39**, 1655–1663 (2014).
14. Handa, T., Harukuni, R. & Fukai, T. Concomitant processing of choice and outcome in frontal corticostriatal ensembles correlates with performance of rats. *Cereb. Cortex* (2021).
15. Zhou, S., Masmanidis, S. C. & Buonomano, D. V. Neural sequences as an optimal dynamical regime for the readout of time. *Neuron* **108**, 651–658 (2020).
16. Shannon, C. E. A mathematical theory of communication. *The Bell system technical journal* **27**, 379–423 (1948).
17. Cover, T. M. & Thomas, J. A. *Elements of Information Theory 2nd Edition* (John Wiley & Sons, Inc., New Jersey, USA, 2006).
18. Schreiber, T. Measuring information transfer. *Phys. review letters* **85**, 461 (2000).
19. Gourévitch, B. & Eggermont, J. J. Evaluating information transfer between auditory cortical neurons. *J. neurophysiology* **97**, 2533–2543 (2007).
20. Bossomaier, T., Barnett, L., Harré, M. & Lizier, J. *An Introduction to Transfer Entropy: Information Flow in Complex Systems* (Springer International Publishing, Cham, Switzerland, 2016).

21. Paluš, M., Komárek, V., Hrnčář, Z. & Štěrbová, K. Synchronization as adjustment of information rates: Detection from bivariate time series. *Phys. Rev. E* **63**, 046211 (2001).
22. Gat, I. & Tishby, N. Synergy and redundancy among brain cells of behaving monkeys. In *Advances in neural information processing systems*, 111–117 (1999).
23. Brenner, N., Strong, S. P., Koberle, R., Bialek, W. & Steveninck, R. R. d. R. v. Synergy in a neural code. *Neural computation* **12**, 1531–1552 (2000).
24. Yeung, R. W. A new outlook on shannon's information measures. *IEEE transactions on information theory* **37**, 466–474 (1991).
25. Schneidman, E. *et al.* Synergy from silence in a combinatorial neural code. *J. Neurosci.* **31**, 15732–15741 (2011).
26. Osborne, L. C., Palmer, S. E., Lisberger, S. G. & Bialek, W. The neural basis for combinatorial coding in a cortical population response. *J. Neurosci.* **28**, 13522–13531 (2008).
27. Emmons, E. B. *et al.* Rodent medial frontal control of temporal processing in the dorsomedial striatum. *J. Neurosci.* **37**, 8718–8733 (2017).
28. Emmons, E. B., Kennedy, M., Kim, Y. & Narayanan, N. S. Corticostriatal stimulation compensates for medial frontal inactivation during interval timing. *Sci. reports* **9**, 1–9 (2019).
29. Oyama, K., Hernádi, I., Iijima, T. & Tsutsui, K.-I. Reward prediction error coding in dorsal striatal neurons. *J. neuroscience* **30**, 11447–11457 (2010).
30. Stalnaker, T. A., Calhoun, G. G., Ogawa, M., Roesch, M. R. & Schoenbaum, G. Reward prediction error signaling in posterior dorsomedial striatum is action specific. *J. Neurosci.* **32**, 10296–10305 (2012).
31. Cools, R. Dopaminergic control of the striatum for high-level cognition. *Curr. opinion neurobiology* **21**, 402–407 (2011).
32. Morita, K., Morishima, M., Sakai, K. & Kawaguchi, Y. Reinforcement learning: computing the temporal difference of values via distinct corticostriatal pathways. *Trends neurosciences* **35**, 457–467 (2012).
33. Yu, L. & Yu, Y. Energy-efficient neural information processing in individual neurons and neuronal networks. *J. Neurosci. Res.* **95**, 2253–2266 (2017).
34. Harris, J. J., Jolivet, R. & Attwell, D. Synaptic energy use and supply. *Neuron* **75**, 762–777 (2012).
35. Harris, J. J., Jolivet, R., Engl, E. & Attwell, D. Energy-efficient information transfer by visual pathway synapses. *Curr. Biol.* **25**, 3151–3160 (2015).
36. Harris, J. J., Engl, E., Attwell, D. & Jolivet, R. B. Energy-efficient information transfer at thalamocortical synapses. *PLoS computational biology* **15**, e1007226 (2019).
37. Rao, R. P. & Ballard, D. H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. neuroscience* **2**, 79–87 (1999).

38. Friston, K. A theory of cortical responses. *Philos. transactions Royal Soc. B: Biol. sciences* **360**, 815–836 (2005).
39. Friston, K. The free-energy principle: a unified brain theory? *Nat. reviews neuroscience* **11**, 127–138 (2010).
40. Keller, G. B. & Mscis-Flogel, T. D. Predictive processing: a canonical cortical computation. *Neuron* **100**, 424–435 (2018).
41. DeLong, M. R. Primate models of movement disorders of basal ganglia origin. *Trends neurosciences* **13**, 281–285 (1990).
42. Ding, J. B., Guzman, J. N., Peterson, J. D., Goldberg, J. A. & Surmeier, D. J. Thalamic gating of corticostriatal signaling by cholinergic interneurons. *Neuron* **67**, 294–307 (2010).
43. Tecuapetla, F., Jin, X., Lima, S. Q. & Costa, R. M. Complementary contributions of striatal projection pathways to action initiation and execution. *Cell* **166**, 703–715 (2016).
44. Klaus, A., Alves da Silva, J. & Costa, R. M. What, if, and when to move: basal ganglia circuits and self-paced action initiation. *Annu. review neuroscience* **42**, 459–483 (2019).
45. Isomura, Y., Harukuni, R., Takekawa, T., Aizawa, H. & Fukai, T. Microcircuitry coordination of cortical motor information in self-initiation of voluntary movements. *Nat. neuroscience* **12**, 1586 (2009).
46. Timme, N. M. & Lapsch, C. A tutorial for information theory in neuroscience. *eneuro* **5** (2018).
47. Timme, N. M., Linsensardt, D. & Lapsch, C. C. A method to present and analyze ensembles of information sources. *Entropy* **22**, 580 (2020).
48. Wibral, M., Vicente, R. & Lizier, J. T. *Directed information measures in neuroscience* (Springer, Berlin, Germany, 2014).
49. Bossomaier, T., Barnett, L., Harré, M. & Lizier, J. *An Introduction to Transfer Entropy: Information Flow in Complex Systems* (Springer, Cham, Switzerland, 2016).

Acknowledgements

The authors would like to gratefully thank Thomas Burns, Milena M. Carvalho, Ruxandra Cojocar, and Yukiko Goda for their assistance and discussions. This research was made possible by the Japan Society for the Promotion of Science (JSPS) research grant (KAKENHI Grant Number: 20K07716 to TH), Hiroshima University Start-up Grant, and RIKEN grant. Illustrations herein were primarily created with BioRender.com

Author contributions statement

T.F. and I.A. designed the project. T.H. performed both the behavioral and electrophysiological experiments, and I.A. analysed the neuronal data. All authors participated in writing the manuscript.

Additional information

Supplementary information includes three figures.

Competing financial interests: The authors declare no competing financial interests.