



CrossMark
click for updates

Opinion piece

Cite this article: Lisman J. 2014 Two-phase model of the basal ganglia: implications for discontinuous control of the motor system. *Phil. Trans. R. Soc. B* **369**: 20130489. <http://dx.doi.org/10.1098/rstb.2013.0489>

One contribution of 18 to a Theme Issue 'The principles of goal-directed decision-making: from neural mechanisms to computation and robotics'.

Subject Areas:

neuroscience

Keywords:

efference copy, motor system, control, tremor, striatum, indirect pathway

Author for correspondence:

John Lisman

e-mail: lisman@brandeis.edu

Two-phase model of the basal ganglia: implications for discontinuous control of the motor system

John Lisman

Biology Department and Volen Center for Complex Systems, Brandeis University, 415 South Street-MS 008, Waltham, MA 02454-9110, USA

In this article, I point out that simple one-phase models of the role of the basal ganglia in action selection have a problem. Furthermore, I suggest a solution with major implications for the organization of the action-selection and motor systems. In current models, the striatum evaluates multiple potential actions by adding biases based on previous conditioning. These biases may arise in both the direct (bias for) and indirect (bias against) pathways. Together, these biases influence which action is ultimately chosen. For efficient conditioning to occur, a positive outcome must selectively strengthen the striatal bias for the chosen action (via a dopaminergic mechanism). This is problematic, however, because all potential action choices have influenced firing patterns in striatal cells during the selection process; it is therefore unclear how the synapses that represent the chosen plan could be selectively strengthened. I suggest a simple solution in which the striatum has two functional phases. In the first phase, the basal ganglia provide biases for multiple potential actions (using both the direct and indirect pathways), leading to the choice of a single action in the cortex. In the second phase, an efference copy of the chosen action is sent to the striatum, where it contributes to the establishment of the eligibility trace for that action. This trace, when acted on by subsequent dopaminergic reinforcement, leads to specific strengthening of the bias only for the chosen action. Consistent with this model, recordings show post-choice imposition onto the striatum of signals corresponding to the chosen action. The existence of dual phases of basal ganglia function implies that decisions about action choice are sent to the motor system in a discontinuous manner. This would not be problematic if the motor system also operated discontinuously. I will review evidence suggesting that this is the case, notably that action is organized by approximately 10 Hz oscillations.

1. Introduction

The basal ganglia have been strongly implicated in the processes by which reinforcement leads to synaptic changes that influence subsequent action choices (for reviews, see [1,2]). Synaptic changes occur at the cortical synapses onto the medium spiny neurons (MSN) of the striatum under the influence of a dopamine signal triggered by reinforcement (for reviews, see [3–5]). This scheme for reinforcement has generally been considered a single-phase process in which action choices can be made in a continuous fashion. The purpose of this brief article is to point out a fundamental problem with such schemes and to propose a solution. Specifically, it is suggested that, for proper conditioning, the basal ganglia must alternate between two phases, one in which multiple actions are weighed and contribute to action choice and one in which the single chosen action is processed and interacts with the reinforcement signal. During this second phase, the basal ganglia cannot function in action choice, making the system fundamentally discontinuous. Experimental evidence consistent with such a two-phase model will be discussed. Furthermore, the possible existence of dual phases in the basal ganglia raises the question of

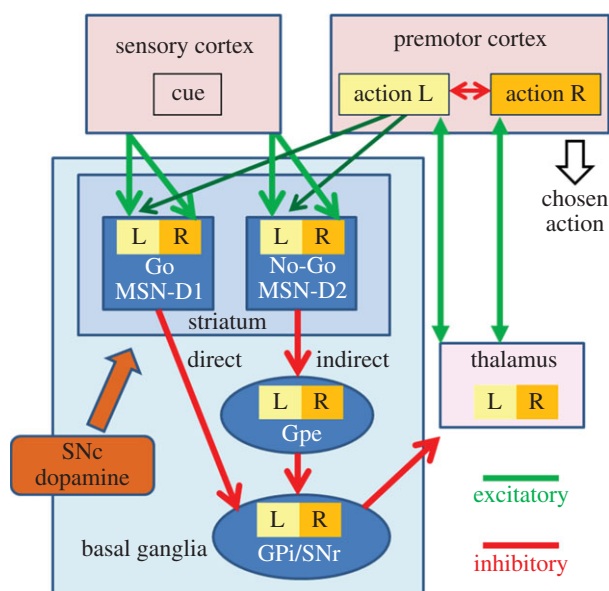


Figure 1. Model of how cortex and basal ganglia lead to action selection during instrumental conditioning. The premotor cortex selects an action (winner-take-all) based on biases for each action received from the basal ganglia. Different potential actions (left/right, L/R) are represented by different groups of medium spiny neurons (MSN) in the striatum in both the direct and indirect pathways. The direct pathway is termed (Go) because activity in this pathway creates a positive bias for an action, whereas the indirect pathway is termed No-Go because it creates a negative bias. These biases are changed by conditioning as follows: conditioning by reward produces an elevation of dopamine and thereby a strengthening of the connections of sensory cortex cells that carry cue information (see below) onto the MSN of the Go (direct) pathway that represent the rewarded action. Negative reinforcement leads to a drop in dopamine, which strengthens connections of sensory cells representing the cue onto MSN of the No-Go pathway. These dopamine-dependent modifications occur only at synapses that carry an eligibility trace. This trace may be determined by previous activity during action selection. Alternatively, as proposed here, the trace is set *after* action selection due to an efference copy of the selected action that is sent from premotor cortex to the striatum; e.g. the signal for L is carried to L MSN-D1 and L-MSN-D2 cells by the pathways shown in dark green (the corresponding pathways for R are not shown). Gpe, globus pallidus externus; GPi, globus pallidus interna; SNc, substantia nigra pars compacta; SNr, substantia nigra pars reticulata. (Online version in colour.)

whether the motor system, which receives instructions from the basal ganglia, is also organized by dual-phase (oscillatory) processes.

Current anatomically specific models of goal-directed instrumental conditioning [6,7] incorporate the finding that there are two major pathways within the basal ganglia, one for promoting actions (via the direct pathway) and one for inhibiting actions (via the indirect pathway) (figure 1). Within the striatum, the input structure of the basal ganglia, MSNs containing D1 receptors give rise to the direct pathway, whereas MSNs containing D2 receptors give rise to the indirect pathway. Both types of MSNs are ‘action specific’ (i.e. firing of a given cell will promote or inhibit a particular action). Both types of MSNs receive cortical inputs that carry information about sensory information. Positive reinforcement is thought to enhance transmission at particular cortical synapses: those that carried information about the cue to the D1-MSNs that promoted the rewarded action.

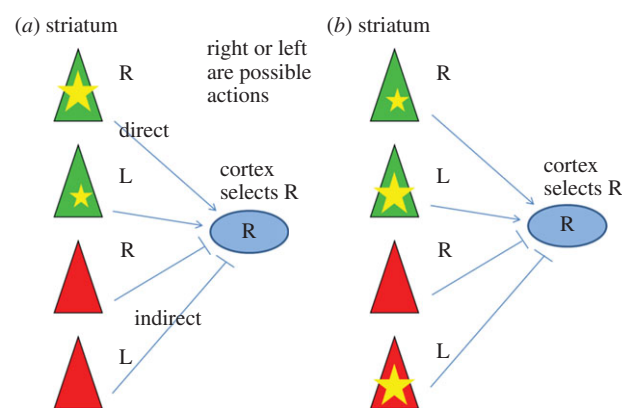


Figure 2. Single-phase models have a problem with correctly assigning credit during reinforcement learning. (a) Activity is only in the MSNs of the direct pathway. Activity levels are designated by the size of the yellow star. Activity in the striatum influences the final action selection by premotor cortex, which in this instance chooses R. In this case, dopamine-dependent reinforcement works correctly to strengthen synapses onto the R MSN in the direct pathway. (b) Both direct and indirect pathways are active in this example, and both ‘vote’ on actions. The final choice is R because votes for L in the direct pathway are balanced by votes against L in the indirect pathway. If R is reinforced, the resulting dopamine release will lead to enhancement of the active synapses onto the L MSNs, which have the largest eligibility trace in the direct pathway. This is problematic because it will not lead to the desired change in behaviour. (Online version in colour.)

Punishment and other forms of negative conditioning enhance transmission onto D2-MSNs and thereby lead to the inhibition of actions. As a result of such changes, the state of synapses onto MSNs comes to depend on the previous history of conditioning. It is envisaged that these synapses ‘vote’ for and against different actions, creating a bias for potential actions in the output structures of the basal ganglia (GPi/SNr). These biases are sent to decision-making structures (e.g. premotor cortex [8]), which, by a winner-take-all process, select the particular action to be executed (but see [9] for a model in which the winner-take-all process is in the basal ganglia itself). A further aspect of current models suggests how positive reinforcement, acting through elevated dopamine, leads to the specific enhancement of the synapses that promoted the selected action. It is envisaged that active synapses set an eligibility trace for synaptic modification. When reward produces a global dopamine signal in the basal ganglia, only synapses with an eligibility trace respond to this dopamine by increasing their strength.

2. A problem with the above model

Figure 2b describes a realistic situation in which this model will not lead to appropriate conditioning. To see the problem, it is useful to first examine the case in which there is only activity in the direct pathway (figure 2a), a situation that does result in appropriate conditioning. Consider a left/right (L/R) choice task in which the experimenter will reward R. In an action selection process leading to the choice of R, the MSNs that stimulated the R action will have fired more vigorously (and thus have a larger eligibility trace) than those stimulating L action. The larger eligibility trace in R MSNs will lead to the

selective strengthening of the cortical synapses onto these cells, thereby producing appropriate conditioning. Now consider the example in figure 2*b*, in which there is activity in both the direct and indirect pathways. This is an important case to consider because experiments show that neurons of both the direct and indirect pathway are active in typical tasks [10]. In the example shown, the activity in the direct pathway is greater for L than for R (leading to a larger eligibility trace in L). However, in the indirect pathway, there is large activity for L; this votes against L in the direct pathway, leading to final selection of R, the desired action. In this example, dopamine will strengthen synapses onto the L cells in the direct pathway because of the large eligibility trace in these cells. This is problematic because it will lead to reduced probability that the desired action (R) will be produced on subsequent trials.

3. Potential solution: a two-phase model

On the one hand, action selection necessarily involves consideration of the pros and cons of various actions; on the other hand, reinforcement would work best if it increased the bias (in the direct pathway) only for the chosen action. A simple way to achieve both requirements is to have two phases of basal ganglia function—in the first, multiple actions are considered; in the second, an efference copy of the single chosen action is imposed on the basal ganglia and specifies the eligibility trace such that only cells representing the chosen action are affected by reward. The setting of the eligibility trace could work as follows. In the first phase, some cortical synapses onto MSNs representing different actions are active, leading to firing of at least some of these cells. These lead to a biochemical mark at active synapses that is a prerequisite for setting the eligibility trace but is not sufficient. In this phase, marks may be set (to various degrees) in cells promoting different actions (e.g. L and R). The second phase occurs after action choice and the arrival of the efference copy at the striatum. If the chosen action is R, only MSNs representing R (considered now only in the direct pathway) will be activated by the efference copy. It may be supposed that the strong activity produced by the efference copy creates the eligibility trace at marked synapses. In summary, active synapses are ‘marked’ and then transformed into an eligibility trace by the efference copy (see [11] for similar ideas about setting the eligibility trace). When the R action is then rewarded, dopamine acts to selectively strengthen synapses in R MSNs (of the direct pathway) that contain an eligibility trace. In summary, this is a three-factor rule: (i) presynaptic activity sets the mark; (ii) the efference copy sets the eligibility trace at marked synapses and (iii) dopamine potentiates synapses that have an eligibility trace. By this mechanism, the resulting synaptic changes will increase the probability that the desired action will be produced on subsequent trials.

The above paragraph focuses on the direct pathway, but what about the indirect pathway? Consider cells in the indirect pathway that represented the chosen action during the choice phase; these were actually voting against the chosen action. If this action is rewarded, it would be appropriate to weaken the cue synapses onto these MSNs of the indirect pathway. There are indications that dopamine has different actions on the direct and indirect pathways and may actually weaken synapses with an eligibility trace [12]. From this perspective, it appears that imposing the efference copy on both

direct and indirect MSNs could work in concert to produce optimal conditioning.

This two-phase solution, while not emphasized in any previous publications, has precedent. In models from the Frank laboratory, a two-phase process is not discussed in the main text but is incorporated into the model, as described in Methods. Indeed, the two phases emerge automatically in an elegant way. In the first phase, connections from cortex to striatum are used to present multiple possible actions to the striatum in parallel, stimulating the basal ganglia to compute biases for these potential actions. These biases are fed back to the cortex and affect the particular action chosen by the cortex. In the second phase, the same connections from cortex to striatum now present only the chosen action to the striatum (i.e. an efference copy). In a similar vein, a model from the Grossberg laboratory [7] is designed so that the action decision (in their case, made in the colliculus) is conveyed to striatum via the thalamus, allowing correct credit assignment. The importance of an efference copy was also introduced by the Fee laboratory as a solution to the credit assignment problem in birdsong learning [13]. That paper raises the possibility that an efference copy might affect the eligibility trace without affecting cell firing, thus potentially solving the credit assignment problem without requiring the two phases of activity proposed here. The models from other groups [9,14] do not incorporate the indirect pathway and thus do not have to deal with credit assignment issues that arise when the indirect pathway is present (figure 2).

4. Evidence consistent with a two-phase model

A critical prediction of the two-phase model is that, before an action choice is made, MSNs will represent several possible actions, whereas after choice, an efference copy will impose activity related only to the chosen action. The study of Lau & Glimcher [15] provides some support for this prediction. The data are complicated because recordings are made from several types of phasically active neurons. However, a general conclusion is that post-choice activity of striatal cells is more driven by the chosen choice than is pre-choice activity. This is particularly evident in two types of cells (directional ‘chosen value cells’ and ‘choice only cells’) that together constitute a substantial fraction of the cells recorded. Figure 4 shows an example. It can be seen that, after action choice ($t = 0$), a brief strong excitation occurs that selectively represents the chosen action. Such activation is likely to result from an efference copy input.

5. Implications of the two-stage (discontinuous) model for motor control

The two-phase model of the basal ganglia implies discontinuous operation of the basal ganglia and raises the question of how such a discontinuous process interfaces with the motor system. Could it be that the motor system also operates discontinuously? Models in which there is a cycling of the cortical regions involved in motor control that matches cycling in the basal ganglia would seem attractive, and one such model is outlined in figure 3. A specific assumption of this model is that the cycle time, during which both basal ganglia phases must occur, is about 100 ms. This is based




Figure 3. Two cycles of two-phase operation of action selection interaction with motor control. The figure illustrates two-phase control of basal ganglia function and how this might interact with the discontinuous (10 Hz) function of the motor system. The relative duration of the two phases is illustrated here as equal but might well be unequal. In this figure, a cycle of BG involvement occurs every 100 ms, illustrating an upper limit. Many tasks may not require BG involvement on each cycle of the motor system. DA, dopamine. (Online version in colour.)

on the growing evidence that the motor system operates discontinuously in a manner organized by approximately 10 Hz oscillations. A full evaluation of this evidence is beyond the scope of this article, but a brief listing of evidence pointing in this direction is summarized in table 1.

According to the figure 3 version of the two-phase model, the basal ganglia could potentially contribute to a different action selection every approximately 10 Hz cycle, a selection that is then executed by the motor system on a cycle-by-cycle basis. Alternatively, perhaps many 10 Hz cycles of the motor system proceed without further involvement of the basal ganglia. There seems to be little evidence relevant to this fundamental design principle. One experiment that examined this issue suggests that both of the above models may be correct under different conditions [30]. Under conditions of unpredicted load (imposition of a viscous drag), there are rapidly occurring corrective actions during a targeting motion, and the basal ganglia blood oxygen level-dependent (BOLD) signal is in proportion to the number of corrective actions. This supports a cycle-by-cycle role of the basal ganglia. However, in a different task in which corrective movements were necessitated by making the target small, basal ganglia BOLD signals were not linked to the number of corrective movements.

6. Conclusion

The question of whether the basal ganglia and motor systems work continuously (one-phase) or discontinuously (two-phase) addresses a fundamental organizational principle of these structures. We argue here for a two-phase model in which one phase operates on multiple possible actions. During this phase, the basal ganglia provide a bias signal for each choice that influences the final choice selection. In the other phase, only the chosen action is represented, a




Figure 4. Recording from a type of striatal cell during choice of saccade direction. The saccade can either be ipsilateral or contralateral to the recording site. This activity after the saccade ($t = 0$) is highly selective for a contralateral choice, consistent with an efference copy signal. Line colour indicates the history of previous reward (value) for the two choices (high, dark blue; low, light yellow). c , period of choice; r , response. (Adapted from [15].) (Online version in colour.)

representation that sets the eligibility trace that makes cells receptive to the dopamine signal. This second phase allows for proper choice evaluation. Consistent with this two-phase model, the results of figure 4 indicate that after action choice, activity corresponding to the chosen action is imposed on a population of striatal neurons for about 100 ms. Additional experiments are needed to understand the role of this subpopulation. Other types of experiments could provide support for the two-phase model. For instance, if there are two phases, there might be a signature of these phases in the field potential. Such a signature is evident for the learning and recall modes of the hippocampus, which can be distinguished by the frequency of gamma oscillations in the CA1 field potential [31–33].

It should be emphasized that there are design principles of action choice and evaluation that need to be addressed and that are glossed over in the model presented here. For instance, the properties and mechanism of eligibility trace are not known

Table 1. Evidence for an oscillatory basis of motor control.

(1)	experiments have shown that single reaching movements have pulses of acceleration at approximately 10 Hz, indicating that single movements contain submovements [16]. Submovements can also be observed by direct recording from muscle; Brown & Cooke [17] found that electromyogram bursts do not have a continuous distribution of durations; instead, they have durations that occur in increments of 70–80 ms. A number of models attribute these submovements to intermittent control [18–20]
(2)	Gross <i>et al.</i> [21] found correlations between oscillations in the cerebello-thalamo-cortical loop, as detected using magnetoencephalography (MEG), and 8 Hz muscle submovements, as detected by EMG. Park <i>et al.</i> [22] showed that muscle oscillations at this frequency (tremor) are dependent on the inferior olive. Previous work [23] had indicated that tremor is still seen in deafferented patients and thus is a result of central drive rather than peripheral feedback
(3)	linkage of action initiation to oscillatory phenomenon is demonstrated by two lines of work. First, movement occurs at a preferential phase of baseline tremor (5–10 Hz) [24]. Second, short latency saccades are initiated at a preferential phase of 10 Hz alpha oscillations [25]
(4)	the frequency limit on single-digit movements is approximately 10 Hz, as demonstrated by the upper limit on typing speed at about 10 characters per second (http://10fastfingers.com/typing-test/english/top50)
(5)	single-unit [26] and field potentials [27,28] reveal oscillations in various frequency bands in motor cortex during actions
(6)	the psychological refractory period occurs when an individual has to complete two tasks that are separated by a certain time interval. If this time interval is too short, there is a bottleneck in performance, suggesting that some portion of processing occurs in a discrete, serial, one-after-the-other manner [29]. A model called the basic unit of motor production (BUMP) assumes three stages of processing: sensory analysis, response planning and a response execution period [18,19]. It is further assumed that the response planning stage is discrete and serial, with a duration of 100 ms. With these assumptions, their model provides justification for the intermittent hypothesis and the psychological refractory period

but will have substantial impact on how conditioning occurs (particularly when reward is delayed). Furthermore, there are subdivisions of the basal ganglia that appear to be organized hierarchically and that deal with high-level (e.g. cognitive) versus low-level (e.g. motor) aspects of behaviour [34]. Whether learning in these subdivisions must occur on separate cycles or can be processed in parallel during the same cycle will have to be dealt with in future models. Finally, the

concept of dual phases postulated here begs the question of what the control mechanism for these phases might be.

Acknowledgements. I thank Paul Glimcher, Michael Frank, Josh Brown and Michale Fee for useful discussions and David Redish for comments on the manuscript. I thank Vivekanand Vimal for help with table 1.

Funding statement. This work was supported in part by the National Institute of Mental Health of the NIH under award no. R01MH102841.

References

- Balleine BW, Delgado MR, Hikosaka O. 2007 The role of the dorsal striatum in reward and decision-making. *J. Neurosci.* **27**, 8161–8165. (doi:10.1523/JNEUROSCI.1554-07.2007)
- Graybiel AM. 2008 Habits, rituals, and the evaluative brain. *Annu. Rev. Neurosci.* **31**, 359–387. (doi:10.1146/annurev.neuro.29.051605.112851)
- Glimcher PW. 2011 Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proc. Natl Acad. Sci. USA* **108**(Suppl. 3), 15 647–15 654. (doi:10.1073/pnas.1014269108)
- Schultz W. 2013 Updating dopamine reward signals. *Curr. Opin. Neurobiol.* **23**, 229–238. (doi:10.1016/j.conb.2012.11.012)
- Surmeier DJ, Plotkin J, Shen W. 2009 Dopamine and synaptic plasticity in dorsal striatal circuits controlling action selection. *Curr. Opin. Neurobiol.* **19**, 621–628. (doi:10.1016/j.conb.2009.10.003)
- Frank MJ, Seeberger LC, O'Reilly RC. 2004 By carrot or by stick: cognitive reinforcement learning in Parkinsonism. *Science* **306**, 1940–1943. (doi:10.1126/science.1102941)
- Brown J, Bullock D, Grossberg S. 2004 How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. *Neural Netw.* **17**, 471–510. (doi:10.1016/j.neunet.2003.08.006)
- Cisek P, Kalaska JF. 2010 Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* **33**, 269–298. (doi:10.1146/annurev.neuro.051508.135409)
- Gurney K, Prescott TJ, Redgrave P. 2001 A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biol. Cybern.* **84**, 401–410. (doi:10.1007/PL00007984)
- 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. (doi:10.1038/nature11846)
- Fee MS. 2012 Oculomotor learning revisited: a model of reinforcement learning in the basal ganglia incorporating an efference copy of motor actions. *Front. Neural Circuits* **6**, 38. (doi:10.3389/fncir.2012.00038)
- Shen W, Flajolet M, Greengard P, Surmeier DJ. 2008 Dichotomous dopaminergic control of striatal synaptic plasticity. *Science* **321**, 848–851. (doi:10.1126/science.1160575)
- Fee MS, Goldberg JH. 2011 A hypothesis for basal ganglia-dependent reinforcement learning in the songbird. *Neuroscience* **198**, 152–170. (doi:10.1016/j.neuroscience.2011.09.069)
- Vitay J, Hamker FH. 2010 A computational model of basal ganglia and its role in memory retrieval in rewarded visual memory tasks. *Front. Comput. Neurosci.* **4**, 13. (doi:10.3389/fncom.2010.00013)
- Lau B, Glimcher PW. 2008 Value representations in the primate striatum during matching behavior. *Neuron* **58**, 451–463. (doi:10.1016/j.neuron.2008.02.021)
- Vallbo AB, Wessberg J. 1993 Organization of motor output in slow finger movements in man. *J. Physiol.* **469**, 673–691.

17. Brown SH, Cooke JD. 1984 Initial agonist burst duration depends on movement amplitude. *Exp. Brain Res.* **55**, 523–527. (doi:10.1007/BF00235283)
18. Bye RT, Neilson PD. 2008 The BUMP model of response planning: variable horizon predictive control accounts for the speed-accuracy tradeoffs and velocity profiles of aimed movement. *Hum. Mov. Sci.* **27**, 771–798. (doi:10.1016/j.humov.2008.04.003)
19. Bye RT, Neilson PD. 2010 The BUMP model of response planning: intermittent predictive control accounts for 10 Hz physiological tremor. *Hum. Mov. Sci.* **29**, 713–736. (doi:10.1016/j.humov.2010.01.006)
20. Crossman ER, Goodeve PJ. 1983 Feedback control of hand-movement and Fitts' Law. *Q. J. Exp. Psychol. A Hum. Exp. Psychol.* **35**, 251–278.
21. Gross J, Timmermann L, Kujala J, Dirks M, Schmitz F, Salmelin R, Schnitzler A. 2002 The neural basis of intermittent motor control in humans. *Proc. Natl Acad. Sci. USA* **99**, 2299–2302. (doi:10.1073/pnas.032682099)
22. Park YG *et al.* 2010 Ca(V)_{3.1} is a tremor rhythm pacemaker in the inferior olive. *Proc. Natl Acad. Sci. USA* **107**, 10 731–10 736. (doi:10.1073/pnas.1002995107)
23. Marsden CD, Meadows JC, Lange GW, Watson RS. 1967 Effect of deafferentation on human physiological tremor. *Lancet* **2**, 700–702. (doi:10.1016/S0140-6736(67)90977-4)
24. Goodman D, Kelso JA. 1983 Exploring the functional significance of physiological tremor: a biospectroscopic approach. *Exp. Brain Res.* **49**, 419–431. (doi:10.1007/BF00238783)
25. Drewes J, VanRullen R. 2011 This is the rhythm of your eyes: the phase of ongoing electroencephalogram oscillations modulates saccadic reaction time. *J. Neurosci.* **31**, 4698–4708. (doi:10.1523/JNEUROSCI.4795-10.2011)
26. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV. 2012 Neural population dynamics during reaching. *Nature* **487**, 51–56. (doi:10.1038/nature11129)
27. Donoghue JP, Sanes JN, Hatsopoulos NG, Gaal G. 1998 Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J. Neurophysiol.* **79**, 159–173.
28. van Wijk BC, Beek PJ, Daffertshofer A. 2012 Neural synchrony within the motor system: what have we learned so far? *Front. Hum. Neurosci.* **6**, 252. (doi:10.3389/fnhum.2012.00252)
29. Pashler H. 1992 Attentional limitations in doing two tasks at the same time. *Curr. Dir. Psychol. Sci.* **1**, 44–48. (doi:10.1111/1467-8721.ep11509734)
30. Grafton S, Tunik E. 2011 Human basal ganglia and the dynamic control of force during on-line corrections. *J. Neurosci.* **31**, 1600–1605. (doi:10.1523/JNEUROSCI.3301-10.2011)
31. Colgin LL, Denninger T, Fyhn M, Hafting T, Bonnevie T, Jensen O, Moser M-B, Moser EI. 2009 Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* **462**, 353–357. (doi:10.1038/nature08573)
32. De Almeida L, Idiart M, Villavicencio A, Lisman J. 2012 Alternating predictive and short-term memory modes of entorhinal grid cells. *Hippocampus* **22**, 1647–1651. (doi:10.1002/hipo.22030)
33. Igarashi KM, Lu L, Colgin LL, Moser MB, Moser EI. 2014 Coordination of entorhinal–hippocampal ensemble activity during associative learning. *Nature* **510**, 143–147. (doi:10.1038/nature13162)
34. Yin HH, Knowlton BJ. 2006 The role of the basal ganglia in habit formation. *Nat. Rev. Neurosci.* **7**, 464–476. (doi:10.1038/nrn1919)