

hypothesis. *Can. Psychiatr. Assoc. J.* **21**, 87–90 (1976).

12. Gordon, W. C. & Spear, N. E. The effects of strychnine on recently acquired and reactivated passive avoidance memories. *Physiol. Behav.* **10**, 1071–1075 (1973).

13. Nader, K., Schafe, G. E. & Le Douarin, J. E. Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval. *Nature* **406**, 722–726 (2000).

14. Debiec, J., LeDoux, J. E. & Nader, K. Cellular and systems reconsolidation in the hippocampus. *Neuron* **36**, 527–538 (2002).

15. Milekic, M. H. & Alberini, C. M. Temporally graded requirement for protein synthesis following memory reactivation. *Neuron* **36**, 521–525 (2002).

16. Fischer, S., Hallschmid, M., Elsner, A. L. & Born, J. Sleep forms memory for finger skills. *Proc. Natl Acad. Sci. USA* **99**, 11987–11991 (2002).

17. Goedert, K. M. & Willingham, D. B. Patterns of interference in sequence learning and prism adaptation inconsistent with the consolidation hypothesis. *Learn. Mem.* **9**, 279–292 (2002).

18. Pace-Schott, E. F. & Hobson, J. A. The neurobiology of sleep: Genetics, cellular physiology and subcortical networks. *Nature Rev. Neurosci.* **3**, 591–605 (2002).

19. Hobson, J. A. & Pace-Schott, E. F. The cognitive neuroscience of sleep: Neuronal systems, consciousness and learning. *Nature Rev. Neurosci.* **3**, 679–693 (2002).

Acknowledgements We thank B. Kocsis, E. Pace-Schott and R. Fosse for comments regarding the study. This research was supported by the National Science Foundation and the National Institute of Mental Health.

Competing interests statement The authors declare that they have no competing financial interests.

Correspondence and requests for materials should be addressed to M.P.W. (mwalker@hms.harvard.edu).

Environmentally mediated synergy between perception and behaviour in mobile robots

Paul F. M. J. Verschure, Thomas Voegtlin* & Rodney J. Douglas

Institute of Neuroinformatics, University/Swiss Federal Institute of Technology (ETH) Zürich, CH-8057 Zürich, Switzerland

* Present address: Institute for Theoretical Biology, Humboldt-University Berlin, D-10115 Berlin, Germany

The notion that behaviour influences perception seems self-evident, but the mechanism of their interaction is not known. Perception and behaviour are usually considered to be separate processes. In this view, perceptual learning constructs compact representations of sensory events, reflecting their statistical properties^{1,2}, independently of behavioural relevance^{3,4}. Behavioural learning^{5,6}, however, forms associations between perception and action, organized by reinforcement^{7,8}, without regard for the construction of perception. It is generally assumed that the interaction between these two processes is internal to the agent, and can be explained solely in terms of the neuronal substrate⁹. Here we show, instead, that perception and behaviour can interact synergistically via the environment. Using simulated and real mobile robots, we demonstrate that perceptual learning directly supports behavioural learning and so promotes a progressive structuring of behaviour. This structuring leads to a systematic bias in input sampling, which directly affects the organization of the perceptual system. This external, environmentally mediated feedback matches the perceptual system to the emerging behavioural structure, so that the behaviour is stabilized.

One reason for the lack of progress in understanding the inter-relationship of behaviour and perception is experimental intractability. An explanation of their coupling requires detailed analysis at both the behavioural and neuronal levels. Our approach was to bypass the animal experimental difficulty by using a mobile robot, for which it is possible to fully observe and quantify perception and

behaviour. The robot is controlled by a neural model, called distributed adaptive control (DAC), that includes mechanisms for perceptual and behavioural learning^{10,11}. The DAC architecture (see Fig. 1 and Methods) consists of three layers: 'reactive', 'adaptive' and 'contextual' control.

The reactive control layer implements a repertoire of basic reflex actions where low-complexity sensory events, unconditioned stimuli (US), trigger simple actions, unconditioned responses (UR), via an internal state (IS) representation. As a result of learning at the level of adaptive control, the purely reactive activation of the IS populations by US events is progressively replaced by acquired representations of sensory events, conditioned stimuli (CS), and the generation of conditioned responses (CR)^{11,12}. US events are the initial reinforcers of this learning process. The local learning mechanism that is used automatically generates a measure, D , of the discrepancy between expected and actual CS events (see Methods). When D falls below a specified transition threshold, θ_D , the contextual control layer is enabled. This layer is a behavioural learning system that constructs higher-order representations of the temporal order of the sensori-motor representations constructed by the adaptive layer (see Methods).

Representations of CS and CR events are stored in short-term memory (STM) when the adaptive layer triggers CRs. The content of STM is stored in long-term memory (LTM) when a goal state is reached, such as when a target is found. CS representations of the LTM of the contextual layer are matched to those generated by the adaptive layer. The best-matching CS representation at the level of contextual control will define the next action by projecting its CR representation onto the motor population M when the reactive layer is quiescent. Chaining through a LTM sequence is achieved through a biased competition mechanism (see Methods). DAC is a practical model of how different learning systems in the mammalian brain act together to generate adaptive goal-oriented behaviour, and is a standard in the field of new artificial intelligence and behaviour-based robotics^{13–16}. Moreover, it exhibits the regularities of bayesian decision-making that are thought to be one of the characteristics of human cognition^{17,18}.

We first investigated the hypothesis that the performance of the robot is enhanced through the contextual layer. To test this hypothesis, we used both simulated and real-world robots in a foraging task where collisions had to be minimized while the number of targets found had to be maximized. We distinguished




Figure 1 Distributed adaptive control. DAC is based on the assumption that adaptive behaviour results from three tightly coupled layers of control: reactive, adaptive and contextual control. Each box represents a neuronal population. Arrows indicate the connections between these populations. US, unconditioned stimulus population. CS, conditioned stimulus population. IS, internal state populations. M, motor neuron population. UR, unconditioned response. CR, conditioned response. STM, short-term memory. LTM, long-term memory. See text and Methods for explanation.

two conditions, in which the contextual layer was either ‘enabled’ or ‘disabled’.

In simulation experiments (see Fig. 2a and Supplementary Information), we found that the adaptive layer improved the performance of the robot through a learning-dependent avoidance of collisions, reflected in the increase of the target/collision ratio (Fig. 2b). We also observed that at the onset of the second stimulation period, the performance of the two conditions diverges: in the enabled condition, performance is strongly enhanced compared with the disabled condition. This difference is due to the activation of the contextual control layer in the enabled condition, as can be deduced from the evolution of the discrepancy measure D (Fig. 2c)—shortly after the onset of the second stimulation period, the D value falls below the transition threshold θ_D .




Figure 2 Experimental protocol and performance in simulated robot experiments for the disabled and enabled conditions using 1,000 exemplars per condition. The initial position and orientation of the robots were randomized. **a**, Each experiment consists of two cycles. Each cycle commences with a stimulation period (2,000 time steps), in which the targets emit a signal (US^+), followed by a recall period (5,000 time steps), in which the target signals are absent. **b**, The ratio between the number of targets found and accumulated collisions averaged in non-overlapping time windows of 100 time steps. **c**, Evolution of the discrepancy measured, D . The dotted horizontal line indicates the value of the transition threshold, θ_D .

During the second stimulation and recall periods, the D value of the enabled condition is markedly below that of the disabled condition. This reduction is accompanied by a significantly lower value of the average absolute change in synaptic efficacies of the connections between the CS and IS populations: enabled $\|\Delta W\| = 2.4 \times 10^{-2}$, disabled $\|\Delta W\| = 2.6 \times 10^{-2}$ (t -test, $P \ll 0.001$). Hence, the transition to contextual control leads to a reduction of the discrepancy between predicted and actual CS events and to a stabilization of the synaptic weights of the adaptive control layer. However, our model has no internal feedback from the contextual to the adaptive control layer: D and ΔW are properties local to the perceptual learning system. Therefore, this difference must be due to the difference in the overt behaviour generated in the two conditions and the systematic bias in the sampling of CS events that this difference causes, that is, behavioural feedback. That is, behaviour is less variable when the contextual layer is enabled, thereby reducing the variability of the sampled sensory inputs. We tested this hypothesis by comparing the entropies of behaviour and sampled stimuli between the two conditions (see Supplementary Information).

We characterized the behavioural entropy, H_B , of the distribution of positions visited for both conditions in an experiment of 10^6 time steps, using the same protocol as above (Fig. 2a). In the disabled condition, H_B was 15.1, whereas the enabled condition showed a lower H_B value of 14.2. These numbers can be compared with the maximal entropy of 15.4, obtained from a uniform distribution of positions; and to an H_B of 11.2 for a minimal cyclic trajectory that follows the shortest path between subsequent targets. The difference in H_B between a uniform distribution of positions and the disabled condition can be explained by the learned avoidance behaviour that causes the robot to avoid the regions close to obstacles. The additional reduction of H_B in the enabled condition is due to a further task-dependent structuring of behaviour (see also real robot results below).

We assessed whether the structuring of behaviour quantified by H_B is associated with a change in the input statistics by calculating the sampling entropy, H_S , of the states of the CS-related sensor. H_S was lower (6.80) for the enabled condition than for disabled (7.95), demonstrating that the structuring of behaviour displayed in the enabled condition is associated with a marked reduction in the variability of the sampled CS events. We deduce that the reduction in behavioural variability, which results from behavioural learning, reduces the set of inputs that the perceptual learning system must classify. As a result of this behavioural feedback, the structures for perceptual learning become adjusted to a smaller set of input states. This change is reflected in the reduction and stabilization of D (Fig. 2c) and the reduced value of the synaptic changes we observed in the enabled condition compared with disabled.

So far we have demonstrated the effect of behavioural feedback on perceptual learning. However, its implications extend beyond sensory classification alone. In our model, the transition to the use of contextual control occurs when D falls below a fixed threshold, θ_D . Hence, behavioural feedback, because it reduces D , may favour the transition to contextual control. To test this hypothesis, we recorded the downward crossings through θ_D for the experiments reported in Fig. 2. On average, these transitions occurred near the onset of the second stimulation period (Fig. 2c). However, in individual experiments, D does not decrease monotonically and small fluctuations of D around θ_D , in the enabled condition, result in oscillations between activating or deactivating contextual control. For the disabled condition, we observed that, on average, oscillations around the transition threshold occurred 6.31 times, whereas for the enabled condition these oscillations were strongly reduced to 3.86. Because there is no mechanism in our model that stabilizes the transition to contextual control, we conclude that this stabilization occurs through behavioural feedback. The switch to contextual control indirectly reduces D through behaviour, and this




Figure 3 Micro-robot experiments. **a**, The CS is provided by a colour charge-coupled device (CCD) camera (1), whereas aversive and appetitive USs, defined as collisions and the level of ambient light, respectively, are provided by the six frontal infrared sensors (2). Locomotion is provided by two wheels (3). **b, c**, The environment is surrounded by a yellow

wall and coloured patches are placed on the floor. The targets (coloured circles) correspond to light sources. The trajectories displayed are those generated by the robot during the recall period, for conditions disabled (**b**) and enabled (**c**). Red, one light source. Green, two light sources. Blue, three light sources.

reduction of D in turn directly favours contextual control. Hence, our simulation experiments show that behavioural feedback not only matches the perceptual system to the emerging behavioural structure, but also in turn stabilizes behavioural learning.

We assessed the generality of our results by applying our model to a similar foraging task using the real-world robot Khepera¹⁹ (Fig. 3a and Supplementary Information). The task requires the robot to learn to avoid collisions (US^-) while visiting illuminated regions in the environment (US^+) as often as possible. The CSs are defined by different colours present in the environment, that is, a yellow wall and coloured patches on the floor.

The trajectories during the recall test show an organization of behaviour consistent with the behavioural entropy measures obtained in the simulation experiments (Fig. 3b, c). In the disabled condition (Fig. 3b), the robot associated yellow patches with aversive events and learned to avoid collisions with the walls of the arena. However, the overall behaviour of the robot consists of following the wall interspersed with deflections off the wall. Thus, the target areas are visited only by chance. In the enabled condition (Fig. 3c), the robot used the coloured patches to reliably locate the areas where it had found positive reinforcement in the past, leading to a more structured behaviour. We quantified this structure by recasting the behaviour as a Markov process. We used an environment with two targets (see Fig. 3b, c, green circles) and estimated the probabilities of the transition between pairs of CS events excluding the yellow wall (Fig. 4 and Supplementary Information).

The disabled condition shows probabilities below 0.25 in transitions between 10 CS pairs. Most of these transitions occur in the periphery of the environment. In the enabled condition, the maximum transition probability increases to 0.75 for a total of 20

CS pairs, including colour patches in the centre of the environment. This pattern is consistent with the trajectories observed earlier (Fig. 3b, c). Moreover, the variability of $\|\Delta W\|$ was much smaller in the enabled compared with the disabled condition consistent with the values found in the simulation experiments: enabled $\|\Delta W\| = 3 \times 10^{-3}$, disabled $\|\Delta W\| = 8 \times 10^{-2}$. This confirms that our results fully generalize to the real world and are robust with respect to the details of the sensory and motor systems employed.

To directly evaluate the effect of behavioural feedback on performance, we compared the enabled condition with a separate control condition, called 'static'. In the static condition, the synaptic efficacies of the adaptive layer were fixed and initialized with the values of those of the enabled condition when the latter stored its first LTM sequence. The static condition can still store STM and LTM sequences but its perceptual learning system is switched off. We found, using both simulated and real-world robots, that performance in the static condition was strongly reduced by comparison with the enabled condition. Moreover, its behavioural trajectory was more variable, its D value higher and the oscillations around θ_D enhanced (see Supplementary Information for details). This result confirms that behavioural feedback directly enhances performance.

Although the microscopic and local view of behaviour and perception has been questioned for being too restricted^{20,21}, it is unclear which other factors should be included to comprehend fully how they are shaped through experience. The robot experiments reported here demonstrate that learning-dependent changes in behaviour can establish a macroscopic feedback loop. Once the contextual layer is activated, the robot can retain and recall




Figure 4 Quantification of the emerging behavioural structure using a hidden Markov model. The transition probabilities between pairs of CS events are mapped on the environment for conditions disabled (**a**) and enabled (**b**). Probabilities exceeding 0.1 are represented by a line that connects the corresponding colour patches. Line thickness reflects the value of the transition probability. Transitions that include the yellow surrounding wall are omitted. The maximum value of the transition probability is 0.25 for the disabled condition and 0.75 for the enabled condition.

successful trajectories, and it develops a habit of following mainly these established paths. In this more restricted and therefore less variable environment, the task of the perceptual learning system is facilitated and biased towards behaviourally relevant sensory events. This improved sensory classification in turn stabilizes the emerging behavioural patterns acquired by the behavioural learning system. Hence, non-neuronal feedback can directly contribute to the organization of behaviour and perception, establishing a synergistic interaction.

To what extent are our robotic results relevant to the understanding of the brain? One prediction of our study is that a change in the afferent input to sensory areas, due to behavioural feedback, can systematically change the organization of perceptual systems. This prediction is supported by the observation that, during development, the organization and response properties of primary sensory areas can be strongly influenced by their afferent inputs²², indicating that during these stages of development behavioural feedback can have a direct impact on the organization of perceptual systems. We expect that behavioural feedback would affect perceptual systems and those structures that readout from these systems,

both during early development and beyond. The observation that the place fields of the CA1 region of the hippocampus show an experience-dependent expansion and shift in their centre of mass in relation to the behavioural trajectory supports this suggestion²³. Moreover, the observation that the responses in the primary auditory cortex are attenuated during vocalization suggests that the brain actively regulates the impact of behavioural feedback²⁴. Thus, we propose that the brain exploits behavioural feedback to constrain perceptual learning and to stabilize acquired behavioural structures.

The role of feedback internal to the neuronal systems underlying perceptual and behavioural learning is receiving increasing attention, particularly in the context of prediction errors of stimulus properties and rewards^{25,26}. Behavioural feedback modifies stimulus sampling and so provides an additional extra-neuronal path for the reduction of prediction errors. Our results provide direct evidence for this assertion by demonstrating that the environmentally mediated feedback between behaviour and perception can significantly affect both of these processes. □

Methods

All neural elements of the DAC architecture are implemented as linear threshold units.

Reactive control

The reactive control layer implements a set of pre-wired reflexes that support basic behaviours where low-complexity sensory events, the US, trigger immediate actions, the UR. The relationship between the US and the UR is established through an intermediate stage that reflects an IS, for example, aversion, ($US^- \rightarrow IS^-$) or appetite ($US^+ \rightarrow IS^+$). US events activate neurons in IS, which in turn activate a population of motor neurons (M), producing an avoidance or approach action (UR), respectively. Action selection is defined by a winner-take-all mechanism in neuronal population M . Conflicts between approach and avoidance actions are resolved by pre-wired relationships between IS populations. If none of the IS populations is active, the reactive control layer generates exploration behaviour, consisting of translational movements.

Adaptive control

The adaptive control layer provides mechanisms for perceptual learning and constructs representations of complex sensory events (CS). These representations arise from the experience-dependent changes in the efficacies of the synaptic connections between the populations of neurons reflecting sensory events (CS) and IS. Activity in IS populations, A_{IS} , is defined by:

$$A_{IS} = A_{US} + WA_{CS} \tag{1}$$

where A_{US} and A_{CS} denote the activity of the US-conveying sensor and CS-driven neural activity, respectively. W represents the synaptic efficacies of the connections between CS and IS populations. The change of these synaptic efficacies, ΔW , is defined by:

$$\Delta W = \eta A_{IS}(A_{CS} - \gamma W^T A_{IS}) \tag{2}$$

where η is a learning rate, γ a linear gain and W^T the transpose of W . Hence, ΔW depends on the difference, or reconstruction error, between the actual CS state, A_{CS} , and that predicted CS state, given the current state of IS and W , $W^T A_{IS}$, also referred to as the sensory expectation, E . This learning rule is related to filtering theory and state estimation²⁷, and similar approaches have been applied to modelling the cortical mechanisms of perceptual learning². Recently, a biophysically realistic implementation of this predictive learning rule was presented²⁸.

D is defined as a leaky average of the difference, $d(A_{CS}, E)$, between the current, A_{CS} , and the predicted CS state, E :

$$D(t+1) = (1 - \alpha_D)D(t) + \alpha_D d(A_{CS}, E)(t) \tag{3}$$

where α_D defines the integration time constant and $d(A_{CS}, E)$ is defined as:

$$d(A_{CS}, E)(t) = \frac{1}{N} \sum_{j=1}^N \left| \frac{A_{CS}(j)}{\max_{1 \leq j' \leq N} A_{CS}(j')} - \frac{E(j)}{\max_{1 \leq j' \leq N} E(j')} \right| \tag{4}$$

where N stands for the size of the CS population.

Contextual control

The contextual control layer provides mechanisms of STM and LTM and is enabled once the discrepancy measure D falls below a fixed threshold, θ_D . This ensures that the representations of CS events, E , used in STM and LTM are based on stable classifications constructed by the perceptual learning system of the adaptive layer. Some key aspects of the contextual layer are: (1) salient CS events, E , and their associated actions, M , are stored in STM, conserving their order of occurrence. Saliency is defined as the interruption of exploration behaviour in the absence of US events. (2) The content of STM is stored in LTM when a goal state is reached, such as finding a target. (3) The sensory content of LTM segments is continuously matched against interpreted sensory events, E , generated by the adaptive layer. (4) LTM segments that match E compete for behavioural control. (5) If the best-matching segment exceeds a specific matching threshold it will control behaviour,

provided that the reactive control layer is inactive. (6) Chaining through LTM sequences is achieved by biasing the LTM matching process.

The competition among LTM segments takes place on a quantity m where m_l^k of segment l of sequence k is defined as:

$$m_l^k = c_l^k t_l^k \quad (5)$$

where c_l^k is defined as $d(E, E_l^k)$ using the definition of d provided in equation (4). $t_l^k \in [0; 1]$, is a dynamic threshold that provides a probabilistic mechanism for chaining through a LTM sequence. In case segment l of sequence k wins the competition by virtue of having the lowest value of m , it will reduce t of segment $l + 1$, of k , t_{l+1}^k , to a fixed value $\beta, \beta \in [0; 1]$. t_l^k relaxes to its default value of 1 according to $t_l^k(t + 1) = \alpha_T + (1 - \alpha_T)t_l^k(t)$, $\alpha_T \in [0; 1]$. The LTM segment that wins the competition will dominate the behavioural output of the overall system when its m is below a fixed threshold.

In the present implementation, STM is a ring buffer with a fixed length of 25 segments, while the capacity of LTM is limited to 64 sequences.

Received 3 March; accepted 7 August 2003; doi:10.1038/nature02024.

1. Olshausen, B. A. & Field, D. J. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**, 607–609 (1996).
2. Rao, R. & Ballard, D. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neurosci.* **2**, 79–87 (1999).
3. Logothetis, N. & Sheinberg, D. Visual object recognition. *Annu. Rev. Neurosci.* **19**, 577–621 (1996).
4. Goldstone, R. Perceptual learning. *Annu. Rev. Psychol.* **49**, 585–612 (1998).
5. Mackintosh, N. *The Psychology of Animal Learning* (Academic, New York, 1974).
6. Lavond, D. G., Kim, J. J. & Thompson, R. F. Mammalian brain substrates of aversive classical conditioning. *Annu. Rev. Psychol.* **44**, 317–342 (1993).
7. Thorndike, E. Animal intelligence: an experimental study of the associative processes in animals. *Psychol. Rev. Ser. Monogr. Suppl.* **2**, 1–109 (1898).
8. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, Massachusetts, 1998).
9. Squire, L. & Kandel, E. *Memory: From Mind to Molecules* (Scientific American Library, New York, 1999).
10. Verschure, P. F. M. J., Kröse, B. & Pfeifer, R. Distributed adaptive control: The self-organization of structured behavior. *Rob. Auton. Syst.* **9**, 181–196 (1992).
11. Verschure, P. F. M. J. & Voegtlin, T. A bottom-up approach towards the acquisition, retention, and expression of sequential representations: Distributed adaptive control III. *Neural Netw.* **11**, 1531–1549 (1998).
12. Verschure, P. F. M. J. & Pfeifer, R. in *From Animals to Animats: Proc. 2nd Int. Conf. Simul. Adapt. Behav. (Honolulu, Hawaii)* (eds Meyer, J. A., Roitblat, H. & Wilson, S.) 210–217 (MIT Press, Cambridge, Massachusetts, 1992).
13. McFarland, D. & Bossert, T. *Intelligent Behavior in Animals and Robots* (MIT Press, Cambridge, Massachusetts, 1993).
14. Clancey, W. *Situated Cognition: On Human Knowledge and Computer Representations* (Cambridge University Press, Cambridge, UK, 1996).
15. Arkin, R. *Behavior-Based Robotics* (MIT Press, Cambridge, Massachusetts, 1998).
16. Pfeifer, R. & Scheier, C. *Understanding Intelligence* (MIT Press, Cambridge, Massachusetts, 1999).
17. Verschure, P. F. M. J. & Althaus, P. A real-world rational agent: Unifying old and new AI. *Cogn. Sci.* **27**, 561–590 (2003).
18. Massaro, D. *Perceiving Talking Faces: From Speech Perception to a Behavioral Principle* (MIT Press, Cambridge, Massachusetts, 1997).
19. Mondada, F., Franzi, E. & Jenne, P. *Experimental Robotics III: Proc. 3rd Int. Symp. Exp. Rob. (Kyoto, Japan, 28–30 October 1993)* 501–513 (Springer, Berlin, 1993).
20. Tolman, E. Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189–208 (1948).
21. Bell, A. Levels and loops: the future of artificial intelligence and neuroscience. *Phil. Trans. R. Soc. Lond. B* **354**, 2013–2020 (1999).
22. Sur, M. & Leamy, C. Development and plasticity of cortical areas and networks. *Nature Rev. Neurosci.* **2**, 251–261 (2001).
23. Mehta, M., Barnes, C. & McNaughton, B. Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc. Natl Acad. Sci. USA* **94**, 8918–8921 (1997).
24. Houde, J., Nagarajan, S., Sekihara, K. & Merzenich, M. Modulation of the auditory cortex during speech: An MEG study. *J. Cogn. Neurosci.* **14**, 1125–1138 (2002).
25. Rescorla, R. & Wagner, A. in *Classical Conditioning 2. Current Theory and Research* (eds Black, A. H. & Prokasy, W. E.) 64–99 (Appleton-Century-Crofts, New York, 1972).
26. Schultz, W. & Dickinson, A. Neuronal coding of prediction errors. *Annu. Rev. Neurosci.* **23**, 473–500 (2000).
27. Kalman, R. A new approach to linear filtering and prediction problems. *Trans. ASME J. Basic Eng.* **82**, 35–45 (1960).
28. Sanchez-Montanes, M., Verschure, P. F. M. J. & König, P. Local and global gating of plasticity. *Neural Comput.* **12**, 519–529 (2000).

Supplementary Information accompanies the paper on www.nature.com/nature.

Acknowledgements We thank A. Baumgartner and J. Manzollini for their support in performing the Markov analysis. Part of this research is supported by the Swiss National Science Foundation, the Volkswagen foundation and the Körber Foundation.

Competing interests statement The authors declare that they have no competing financial interests.

Correspondence and requests for materials should be addressed to P.F.M.J.V. (pfmjv@ini.phys.ethz.ch).

Regulation of neuroblast competence in *Drosophila*

Bret J. Pearson & Chris Q. Doe

Institutes of Neuroscience and Molecular Biology, Howard Hughes Medical Institute, University of Oregon 1254, Eugene, Oregon 97403, USA

Individual neural progenitors generate different cell types in a reproducible order in the retina^{1–3}, cerebral cortex^{4–6} and probably in the spinal cord⁷. It is unknown how neural progenitors change over time to generate different cell types. It has been proposed that progenitors undergo progressive restriction⁸ or transit through distinct competence states^{9,10}; however, the underlying molecular mechanisms remain unclear. Here we investigate neural progenitor competence and temporal identity using an *in vivo* genetic system—*Drosophila* neuroblasts—where the Hunchback transcription factor is necessary and sufficient to specify early-born cell types¹¹. We show that neuroblasts gradually lose competence to generate early-born fates in response to Hunchback, similar to progressive restriction models⁸, and that competence to acquire early-born fates is present in mitotic precursors but is lost in post-mitotic neurons. These results match those observed in vertebrate systems, and establish *Drosophila* neuroblasts as a model system for the molecular genetic analysis of neural progenitor competence and plasticity.

Despite substantial progress in vertebrates, we still know little about the molecular basis for how a single neural precursor sequentially generates different cell types. This is primarily due to the lack of an *in vivo* model system where a single neural progenitor can be studied at reproducible times during its lineage. The *Drosophila* embryonic central nervous system (CNS) lends itself well to the study of neural progenitor plasticity because neural progenitors (neuroblasts (NBs)) can be individually identified, each NB generates different cell types in a reproducible order, molecular markers exist for each of these cell types, intrinsic factors are known that confer different temporal identities, and gene expression can be readily manipulated at specific points within the NB lineage¹¹. NBs repeatedly divide in a stem-cell-like mode to ‘bud off’ a series of smaller daughter cells called ganglion mother cells (GMCs). Cell lineage studies show that every GMC has a unique identity based on its ‘birth’ order within the NB lineage, and generates a characteristic pair of neurons or glia. Recently, four transcription factors have been identified that are excellent candidates for specifying GMC temporal identity^{11,12}. NBs sequentially express the transcription factors Hunchback (Hb)→Krüppel→Pdm1→Castor, with GMCs inheriting the transcription-factor profile of the parental NB on their generation, which is then maintained in their own neuronal progeny (Fig. 1)¹¹. *hb* is both necessary and sufficient for specifying the first-born temporal identities in multiple NB lineages, even though first-born cells can be motor neurons, interneurons, or glia¹¹. Here we manipulate the timing and levels of Hb in a model NB lineage (NB7-1) to ask two fundamental questions: when does a NB lose competence to generate first-born neurons (immediately after Hb downregulation, progressively during its lineage, or never)? And when during neuronal differentiation is competence to generate first-born neurons lost (in NBs, GMCs, or post-mitotic neurons)?

To assay changes of temporal identity within a single lineage, we precisely define the birth order and sibling relationships of the early NB7-1 lineage. NB7-1 generates five motor neurons (U1–U5) and about 30 interneurons^{13,14}. The five U motor neurons have stereotyped positions in the CNS (Fig. 1a), express the Even-skipped (Eve) transcription factor (Fig. 2a), and innervate specific body wall