What are the computational benefits of stereotyped connectivity in the lateral horn? An obvious, yet intriguing, hypothesis is that connectivity in the lateral horn has evolved to optimize the representation of odor components that are predictable and relevant for a species. More specifically, it may be hypothesized that LHNs become tuned to ecologically important features of natural odors, much like higher order neurons in the visual system are tuned to informative features of visual scenes. Consistent with this possibility, PNs converging onto type I* LHNs appear to respond to fruity odors, which may be relevant to fruit flies. However, systematic and exhaustive analyses of natural odor space will eventually be necessary to address this question. A more extreme hypothesis is that individual subtypes of LHNs become tuned to odor components that directly control behavioral outputs. This hypothesis views the lateral horn as a switchboard that interfaces complex sensory inputs to defined behavioral outputs. Stimulation or silencing of specific LHNs should therefore elicit or suppress distinct behaviors, a prediction that may be tested by opto- or pharmacogenetic experiments. These hypotheses are not identical, not mutually exclusive and certainly not exhaustive. Indeed, type I and type II LHNs differ not only in their odor selectivity, but they also project to distinct subregions of the protocerebrum. It is therefore possible that subpopulations of LHNs use different mechanisms to integrate PN input to fulfill different functions.

Generally, the stereotyped connectivity between PNs and LHNs reinforces the idea that the lateral horn processes odor information that is predictable on evolutionary timescales, whereas the mushroom body modifies odor processing on the basis of an individual's experience. However, stereotyped connectivity does not rule out the possibility that odor processing in the lateral horn is subject to modification by experience-dependent plasticity. Moreover, the two processing streams appear to converge again in the lateral horn and possibly in other brain areas[10,14]. It will be interesting to examine whether odor-evoked activity of LHNs or downstream neurons can be modified by plasticity processes or by mushroom body outputs.

Similar to the PNs of insects, output neurons of the vertebrate olfactory bulb project to multiple target areas. Projections to the cortical amygdala, a brain region that likely controls defined behaviors, exhibit a coarse topography, whereas projections to piriform cortex, a large associative memory area, appear to lack topography[15]. Assuming that topography indicates stereotyped connectivity, the lateral horn and the cortical amygdala may therefore exhibit not only functional similarities, but also similar connectivity features. However, more information is required about both brain areas to understand how far this similarity extends. In *Drosophila*, the results of Fisek and Wilson[8] mark a fundamental step forward in our understanding of the mechanisms by which olfactory stimuli influence animal behavior.

COMPETING FINANCIAL INTERESTS
The authors declare no competing financial interests.

1. Harris, K.D. & Mrsic-Flogel, T.D. *Nature* **503**, 51–58 (2013).
2. Ganguli, S. & Sompolinsky, H. *Annu. Rev. Neurosci.* **35**, 485–508 (2012).
3. Laurent, G. *Nat. Rev. Neurosci.* **3**, 884–895 (2002).
4. Wiechert, M.T., Judkewitz, B., Riecke, H. & Friedrich, R.W. *Nat. Neurosci.* **13**, 1003–1010 (2010).
5. Murthy, M., Fiete, I. & Laurent, G. *Neuron* **59**, 1009–1023 (2008).
6. Caron, S.J., Ruta, V., Abbott, L.F. & Axel, R. *Nature* **497**, 113–117 (2013).
7. Gruntman, E. & Turner, G.C. *Nat. Neurosci.* **16**, 1821–1829 (2013).
8. Fisek, M. & Wilson, R.I. *Nat. Neurosci.* **17**, 280–288 (2014).
9. Masse, N.Y., Turner, G.C. & Jefferis, G.S. *Curr. Biol.* **19**, R700–R713 (2009).
10. Heisenberg, M. *Nat. Rev. Neurosci.* **4**, 266–275 (2003).
11. Jefferis, G.S. *et al. Cell* **128**, 1187–1203 (2007).
12. Ruta, V. *et al. Nature* **468**, 686–690 (2010).
13. Gupta, N. & Stopfer, M. *J. Neurosci.* **32**, 8138–8148 (2012).
14. Séjourné, J. *et al. Nat. Neurosci.* **14**, 903–910 (2011).
15. Sosulski, D.L., Bloom, M.L., Cutforth, T., Axel, R. & Datta, S.R. *Nature* **472**, 213–216 (2011).

Motor variability is not noise, but grist for the learning mill

David J Herzfeld & Reza Shadmehr

A study demonstrates that variability in how people perform a movement can predict the rate of motor learning on an individual basis. This suggests that motor 'noise' is a central component of motor learning.

As anyone who has tried to learn a new sport can attest, repeatedly performing a movement does not result in the same motor output on every attempt. Rather, there is variability in our motor commands. Until now, it was thought that this variability is due to noise, something that should be avoided. However, a study in this issue of *Nature Neuroscience*[1] demonstrates that variability has a crucial correlate: it can predict the rate of motor learning on an individual basis. That is, variability of movements can predict who will learn a specific motor task faster. These results provide intriguing evidence that some of the motor variability commonly attributed to unwanted noise is in fact exploration in motor command space.

When we make a movement and experience an error, on the next attempt our brain updates the motor commands to compensate for some fraction of the error[2]. Mathematically, this can be written in its simplest form as an equation that relates the participant's motor commands on the nth trial, $u(n)$, to the motor command on trial $n+1$: $u(n+1) = u(n) + \eta e(n)$, where η is the error sensitivity and e is the error experienced. For example, say you are playing basketball and attempt a free throw. Your ball misses the basket to the left, exhibiting error e. The amount that you adjust the motor commands to your arm and shoot the next ball is a reflection of η, your sensitivity to error. This single equation, and its extensions[3,4], account for a vast array of motor learning data. However, the error sensitivity term, η, varies substantially from individual to individual and task to task. What makes some individuals more sensitive to error and therefore faster learners?

Wu *et al.*[1] investigated this question, with the hypothesis that the rate of learning in a novel motor task may be related to the amount of task-related variability, or noise, that each subject naturally expresses in their baseline movements. This motor noise, which until now was thought to be an unwanted feature of movements,

David J. Herzfeld and Reza Shadmehr are in the Department of Biomedical Engineering, Laboratory for Computational Motor Control, Johns Hopkins University School of Medicine, Baltimore, Maryland, USA.
e-mail: shadmehr@jhu.edu

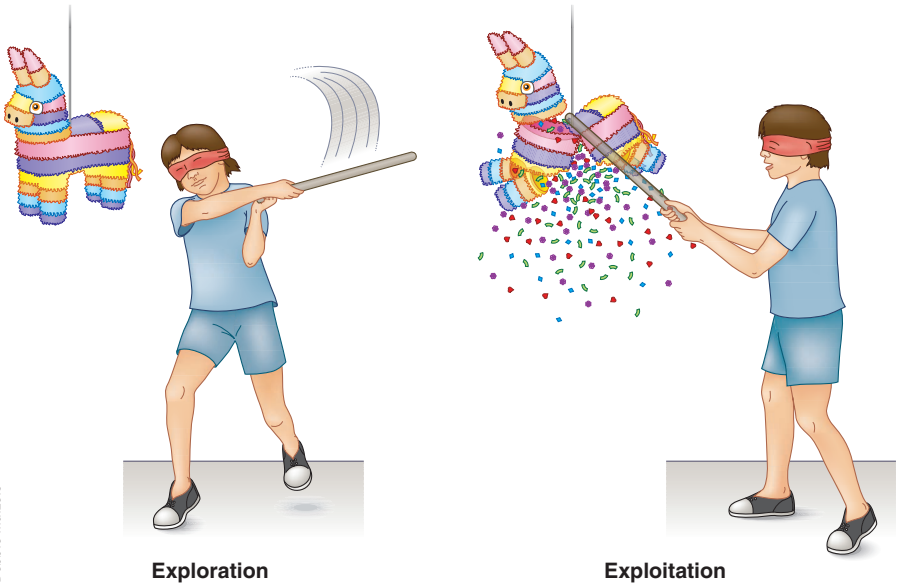

Debbie Maizels

Figure 1 Motor variability. In trying to find the piñata, the child begins by exploring the space of possible swing locations, with a large amount of variability between each stroke. Once contact is made, the child switches to a strategy of exploitation, in which the same movement is repeated on subsequent swings.

something to be avoided and reduced with practice, is shown to have a very useful feature: it allows the brain to learn faster.

Consider a classic children's game, hitting a piñata. In the game, the child is typically blindfolded and told to hit the piñata until breaking it (**Fig. 1**). The child begins the game by swinging wildly all around, attempting to locate the hanging piñata. However, once contact is made, the location where the child swings is similar from one swing to another. These two strategies, exploration and exploitation, are the basis for reinforcement learning. The authors hypothesized that better exploration (characterized by movement variability) in the task-specific dimension may contribute to a faster learning rate (exploitation).

To test their hypothesis, in their first experiment Wu *et al.*[1] considered a reaching task in which participants were asked to move their arm so that it matched a guide trajectory. Subjects could not see their movements, and so they simply tried as best as they could to match the presented trajectory. However, unbeknown to them, the authors were measuring the variability of the movements along a dimension that was irrelevant to the guide trajectory but relevant to an unrelated trajectory, one that their movements would be evaluated against in a subsequent task. If the variability in this baseline task was not simply noise but a form of natural exploration, then people who showed a greater variability during baseline should exhibit faster learning of the subsequent task.

Once the baseline period ended, in the training period subjects received a numerical score after each trial (a reinforcement signal), providing feedback on how well they were able to match the new goal trajectory, although they still could not see their movements. The authors found a significant correlation between the amount of task-relevant variance during the baseline period and the rate of motor learning, supporting the hypothesis that variability in the baseline period was not noise, but exploration, which then predicted faster learning of the task.

Reinforcement learning, which relies on learning from reward prediction error (for example, discrepancy between predicted and observed score of a movement), likely relies on different neural circuitry than learning from sensory prediction error (for example, discrepancy between predicted and observed proprioceptive feedback). The former is thought to rely on the basal ganglia[5], whereas the latter is thought to rely on the cerebellum[6]. Therefore, the correlation that the authors observed in their first experiment may only be true in reinforcement learning tasks and not when learning relies on sensory prediction errors. In another experiment, the authors addressed this by first measuring variability of movements in a baseline reaching task and then using that variability to predict how fast each subject could learn to reach in a force-field motor adaptation task. This form of motor learning is largely dependent on sensory prediction errors and is critically dependent on the cerebellum[7]. The authors found that the variability along a specific dimension during the baseline condition (the dimension relevant to learning the movement) was an excellent predictor of the subsequent learning rate in the force-field task.

If variability contributes to a faster learning rate, then would an increase in the learning rate also result in an increase in task-specific motor variability? To answer this question, the authors designed an experiment in which perturbations were likely to repeat. With subsequent repetitions of the perturbation, participants responded by increasing their learning rate upon re-exposure to the same perturbation (that is, η increased). The authors found that after increases in participants' learning rates, the motor variability increased along a dimension associated with the type of perturbation that was just learned. For instance, when the forces applied to the hand were velocity-dependent, the largest component of the variance measured during baseline conditions became more velocity-dependent after repeated exposure. This implies that modulation of learning rate coincides with modulation of variance along the task-relevant dimension.

Until now, motor variability has been viewed as an unwanted feature of movements, a noise that the brain is able to reduce only with practice[8]. However, Wu *et al.*[1] show that task-relevant motor variability, measured during baseline before people are exposed to a novel motor task, can be used to predict the rate of learning in the task. This relationship exists in two different forms of learning: paradigms that depend on reward prediction errors and paradigms that depend on sensory prediction errors. Therefore, a portion of what has been historically considered motor 'noise' is in fact an asset used by the brain to promote learning.

COMPETING FINANCIAL INTERESTS
The authors declare no competing financial interests.

1. Wu, H., Miyamoto, Y., Gonzales-Castro, L.N., Olveczky, B.P. & Smith, M.A. *Nat. Neurosci.* **17**, 312–321 (2014).
2. Thoroughman, K.A. & Shadmehr, R. *Nature* **407**, 742–747 (2000).
3. Smith, M.A., Ghazizadeh, A. & Shadmehr, R. *PLoS Biol.* **4**, e179 (2006).
4. Kording, K.P., Tenenbaum, J.B. & Shadmehr, R. *Nat. Neurosci.* **10**, 779–786 (2007).
5. Schultz, W., Dayan, P. & Montague, P.R. *Science* **275**, 1593–1599 (1997).
6. Tseng, Y.W., Diedrichsen, J., Krakauer, J.W., Shadmehr, R. & Bastian, A.J. *J. Neurophysiol.* **98**, 54–62 (2007).
7. Smith, M.A. & Shadmehr, R. *J. Neurophysiol.* **93**, 2809–2821 (2005).
8. Shmuelof, L., Krakauer, J.W. & Mazzoni, P. *J. Neurophysiol.* **108**, 578–594 (2012).