In nonhuman primates, DENV infection did not affect the outcome of subsequent ZIKV infection (*5*). However, population studies showed that previous DENV immunity correlated with lower risk of ZIKV infection (*6*). In the Nicaraguan cohort, there was evidence that previous DENV infection may protect against subsequent symptomatic ZIKV infection. However, there was no relationship with DENV immunity, and no enhancing effect was observed (*7*).

Therefore, studies on how DENV infection affects subsequent ZIKV infection outcome have shown varied results. The reasons for this could include differences in the number of prior DENV exposures and serotypes; time between DENV and ZIKV infections; and differences between animal model, in vitro, and human responses. In addition, the range of outcomes for Zika disease is not as wide as for dengue, so differences due to any enhancement could be harder to detect. There are also remaining questions about how prior DENV infection modulates the risk of birth defects to infants born to mothers infected with ZIKV during pregnancy.

What about how ZIKV infection modulates subsequent DENV infection outcome? Before the 2015–2016 outbreak, there had been little research on ZIKV and few reported cases, so the long-term impact of ZIKV infection could not be assessed. Initially after the outbreak, it was possible to show that individuals who were exposed to ZIKV generated both neutralizing and enhancing antibodies for DENV (*8*). In nonhuman primates, there was some evidence of enhancement upon infection with DENV a short time after ZIKV infection (*9*). Now, 5 years after the outbreak, it is possible to assess the longer-term effects of ZIKV on subsequent dengue disease risk in humans. A previous modeling study using reported dengue cases from Colombia and Brazil estimated that ZIKV infection led to cross-protection to DENV, resulting in low reported dengue case numbers in the 2 years after the ZIKV outbreak (*10*). This study predicted future higher numbers of dengue cases after this period, which will need to be assessed.

The study of Katzelnick *et al.* similarly shows very low DENV transmission in the 2 years after the ZIKV outbreak. The important aspect of this study is in the assessment of individual disease risk in the 2019 dengue outbreak, the largest on record. During this epidemic, Katzelnick *et al.* clearly show that on an individual level, there is an increased risk of symptomatic and severe dengue disease in individuals with prior ZIKV immunity. Together, these data show that ZIKV infection modulates the disease severity of subsequent DENV infection in both negative and positive ways, on an individual and population level. This suggests that previous ZIKV infection acts similarly to previous DENV infection. However, there are still unknowns, such as whether risk differs by DENV serotype and the longer-term impact of ZIKV infection on dengue.

Another reason this is an important result is for vaccination. Follow-up of participants in the Dengvaxia vaccine trial showed that dengue-unexposed individuals who were vaccinated showed an increased risk of severe disease in the longer term upon the next DENV infection, similar to individuals experiencing a second natural DENV infection (*11*). Dengvaxia is the only licensed dengue vaccine to date, but there are others in development. For ZIKV, there are also several vaccines currently in development, but none are licensed. It is now important to establish whether dengue vaccine efficacy varies according to previous ZIKV infection. For Dengvaxia, vaccination in the initial trials was before the 2015–2016 ZIKV outbreak. It is also possible that ZIKV vaccination could lead to more severe subsequent DENV infections. It would therefore seem prudent for both DENV and ZIKV vaccine studies to measure immunity to both viruses before and after vaccination and to follow up for both diseases within any trial. Because viral immunity will be changing over time, long-term follow-up will be required.

For other flaviviruses, this interaction with vaccination and risk has also been considered, such as for Japanese encephalitis virus and DENV (*12*). Interactions with the yellow fever flavivirus and vaccinations should also be considered in relevant countries. More generally, the cross-reactive properties of viral immunity, how to measure it, how it changes over time, and the impact on disease risk are important for a number of viral infections, and the study of Katzelnick *et al.* serves as a reminder that issues of cocirculation and cross-immunity should remain in view for understanding transmission dynamics and vaccination. ■

REFERENCES AND NOTES

1. W. Dejnirattisai *et al.*, *Science* **328**, 745 (2010).
2. WHO, Zika virus and complications: 2016 Public Health Emergency of International Concern; www.who.int/emergencies/zika-virus-tmp/en/.
3. L. C. Katzelnick *et al.*, *Science* **369**, 1123 (2020).
4. W. Dejnirattisai *et al.*, *Nat. Immunol.* **17**, 1102 (2016).
5. M. K. McCracken *et al.*, *PLOS Pathog.* **13**, e1006487 (2017).
6. I. Rodriguez-Barraquer *et al.*, *Science* **363**, 607 (2019).
7. A. Gordon *et al.*, *PLOS Med.* **16**, e1002726 (2019).
8. K. Stettler *et al.*, *Science* **353**, 823 (2016).
9. J. George *et al.*, *Sci. Rep.* **7**, 10498 (2017).
10. R. K. Borchering *et al.*, *Nature* **10**, 5730 (2019).
11. S. Sridhar *et al.*, *N. Engl. J. Med.* **379**, 327 (2018).
12. K. B. Anderson *et al.*, *PLOS Negl. Trop. Dis.* **5**, e1311 (2011).

10.1126/science.abd5922

NEUROSCIENCE

Strategies for navigating a dynamic world

A rich neural representation of its environment enables an adaptable organism to respond to changes

By **Saurabh Steixner-Kumar** *and* **Jan Gläscher**

One of the most difficult problems for an adaptable agent is gauging how to behave in a nonstationary environment. When conditions are stable, an organism generally pursues a strategy known to provide the best outcome. However, when environmental conditions change, an organism abandons the current action plan and searches for a new best option. The most challenging aspect of this search—calculating the exact time point at which to change strategies—requires the brain to integrate past and present observations and evaluate whether they remain consistent with current environmental conditions. On page 1076 of this issue, Domenech *et al.* (*1*) report on the modeling of rare direct electrical recordings from the prefrontal cortices (PFCs) of a small group of human epilepsy patients as they flexibly negotiated a nonstationary environment.

To understand the brain's mode of navigation, consider for example a sailor at sea (see the figure). The winds and the currents determine the waves that drive the sailor to continuously adjust the rudder so as to stay on course. By observing the wave patterns, he can anticipate the navigational effects of his actions and adapt accordingly. But when the currents or the weather changes, the sailor must adapt his course to reach the next port of call. At that time, the sailor observes essentially the same stimulus (the waves) but has to remap his action plan (rudder adjustments) to the new wind conditions and currents.

This difficult-decision problem—how to detect and then adapt to a nonstationary environment—is captured perfectly in the exploration-exploitation dilemma: When

Institute for Systems Neuroscience, University Medical Center Hamburg-Eppendorf, Hamburg, Germany.
Email: glaescher@uke.de; s.steixner-kumar@uke.de

should I stop exploiting my current action plan and start exploring different ways to reach my goals? An optimal solution tracks the discounted sum of normalized future rewards. However, this approach applies strictly to stationary environments and thus does not capture the dynamic changes that organisms encounter in their daily lives (*2*). Yet the human brain and those of other species seem to smoothly solve the exploration-exploitation dilemma in nonstationary environments.

Decision neuroscience has investigated the flexible adaptation to changing environmental contingencies with diverse experimental paradigms and assorted computational models. The simplest paradigm is probabilistic reversal learning, in which the agent has to search for reward among two options with complementary reward probabilities. This adaptation problem can be solved by hidden Markov models (*3*), which are well-approximated by reinforcement learning (RL) models that also update nonchosen actions (*4*). Extension of this paradigm to include independently changing reward probabilities reveals two distinct neural responses: Expected-value signals, which reflect "exploitative" choices, spur activation of the ventromedial prefrontal cortex (vmPFC); and "explorative" choices (that is, the choosing of a currently lesser valued option) activate the frontopolar cortex (*5*).

Another task with both rapid and slow changes in the reward probabilities of various options was used to develop a hierarchical Bayesian model that estimates the volatility of the environment and adjusts the learning rate accordingly (*6*). This model has found its generalization in the hierarchical Gaussian filter (HGF) framework (*7*), which is widely used in modeling social and nonsocial human decision-making in nonstationary environments. Although these computational modeling frameworks differ, all are trying to solve similar problems: How to infer the latent structure of the world from discrete observations and how to detect transitions between different states of the world.

Domenech *et al*. address the same problems with yet another experimental paradigm, this one carried out with a small group of human epilepsy patients. Electrodes deeply implanted in the patients' PFCs delivered direct electrical recordings from the vmPFC and dorsomedial PFC (dmPFC) while the patients performed a multioption decision task. The participants had to associate three different stimuli with three distinct actions, thus constituting an action plan. The mapping changed every 33 to 57 trials, and participants had to relearn the association of the same stimuli with a different combination of actions, much like our sailor at sea who faces changes in weather and currents that alter wave patterns. The computational model (*8*) generates a reliability value for the ongoing action plan and other concurrently monitored plans. When the ongoing action plan is deemed reliable, the model is in "exploitation" mode and learns the stimulus-action mapping through RL mechanisms. When the ongoing action plan is deemed unreliable, the model switches to "exploration" mode. New provisional action plans are created and evaluated, until one emerges as a reliable predictor for successful stimulus-action mapping (see the figure).

Using a state-of-the-art model-based analysis that associates the model-derived variables with the brain activity in various frequency bands of the neural recordings, the authors found a delicate interplay between the vmPFC and dmPFC that supports a predictive coding interpretation for resolution of the exploration-exploitation dilemma. vmPFC monitors and represents the reliability of the ongoing action plan. vmPFC relays the ongoing action plan to the dmPFC as either a "stay" or "switch" trial. A stay trial triggers additional learning through RL mechanisms in the dmPFC. In contrast, the dmPFC responds to a switch trial by suppressing activity related to maintaining the ongoing action plan. These findings resonate with and extend earlier results obtained with functional neuroimaging (*5, 9*).

These computational approaches to the problem of behavioral flexibility in a nonstationary environment share one commonality: They are all building a model of the environment and the transition therein, either explicitly (as in the HGF framework) or implicitly (by evaluating the ongoing action plan, as in the Domenech *et al*. study). Although all of these models strive for generality, each was developed for a specific experimental context. It remains to be seen which of these provides the best account of flexible decision-making in humans and other species, preferably using a unified experimental paradigm. A model-free RL account (*10*) likely will not suffice, as several studies have demonstrated the superiority of more-complex models over this "vanilla" RL model. Rather, an agent requires a rich representation of the environment and its dynamic transitions (often referred to as model-based learning) (*10*) to solve the exploration-exploitation dilemma and flexibly respond to a changing world. ■

A sailor solves a dilemma at sea

As the ship nears bad weather, the sailor's ventromedial prefrontal cortex (vmPFC) evaluates the ongoing (orange) action plan (exploitation) and the prospective (brown, red) plans (exploration). Once the red (calm waters) plan is exploited, the sailor's dorsomedial PFC (dmPFC) uses trial-and-error learning to map the proper rudder adjustments.

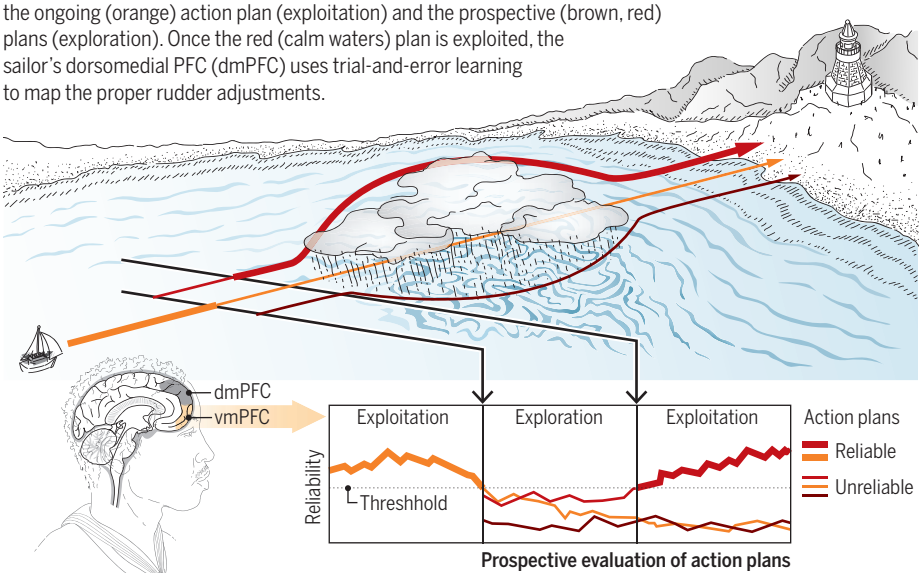

GRAPHIC: A. KITTERMAN/*SCIENCE*

REFERENCES AND NOTES

1. P. Domenech, S. Rheims, E. Koechlin, *Science* **369**, eabb0184 (2020).
2. J. D. Cohen, S. M. McClure, A. J. Yu, *Philos. Trans. R. Soc. London Ser. B* **362**, 933 (2007).
3. A. N. Hampton, P. Bossaerts, J. P. O'Doherty, *J. Neurosci.* **26**, 8360 (2006).
4. J. Gläscher, A. N. Hampton, J. P. O'Doherty, *Cereb. Cortex* **19**, 483 (2009).
5. N. D. Daw, J. P. O'Doherty, P. Dayan, B. Seymour, R. J. Dolan, *Nature* **441**, 876 (2006).
6. T. E. J. Behrens, M. W. Woolrich, M. E. Walton, M. F. S. Rushworth, *Nat. Neurosci.* **10**, 1214 (2007).
7. C. Mathys, J. Daunizeau, K. J. Friston, K. E. Stephan, *Front. Hum. Neurosci.* **5**, 39 (2011).
8. A. Collins, E. Koechlin, *PLOS Biol.* **10**, e1001293 (2012).
9. M. Donoso, A. G. E. Collins, E. Koechlin, *Science* **344**, 1481 (2014).
10. N. D. Daw, P. Dayan, *Philos. Trans. R. Soc. London Ser. B* **369**, 20130478 (2014).

10.1126/science.abd7258

Strategies for navigating a dynamic world

Saurabh Steixner-Kumar and Jan Gläscher

Science **369** (6507), 1056-1057.
DOI: 10.1126/science.abd7258