MIND HACKS™

Tips & Tools for Using Your Brain

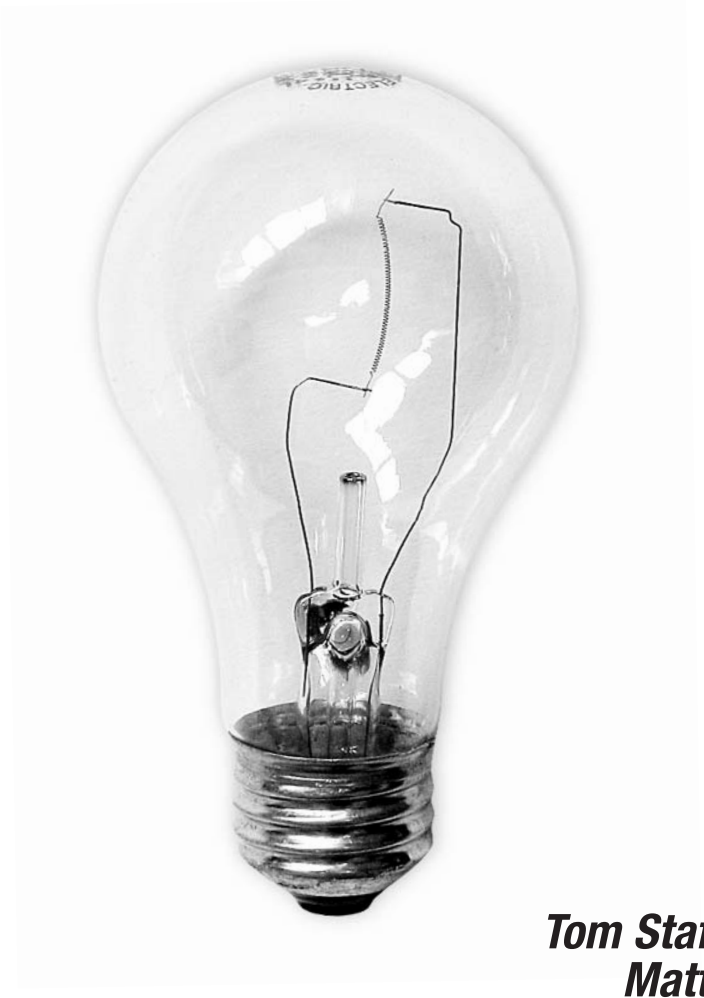

Tom Stafford &
Matt Webb

O'REILLY®

Foreword by Steven Johnson, author of Mind Wide Open

HACK #33 Neural Noise Isn't a Bug; It's a Feature

Neural signals are innately noisy, which might just be a good thing.

Neural signals are always noisy: the timings of when they fire, or even whether they fire at all, is subject to random variation. We make generalizations at the psychological level, such as saying that the speed of response is related to intensity by a certain formula—Pieron's Law **[Hack #11]**. And we also say that cells in the visual cortex respond to different specific motions **[Hack #25]**. But both of these are true only *on average*. For any single cell, or any single test of reaction time, there is variation each time it is measured. Not all the cells in the motion-sensitive parts of the visual cortex will respond to motion, and those that do won't do it exactly the same each time we experience a particular movement.

In the real world, we take averages to make sense of noisy data, and somehow the brain must be doing this too. We know that the brain is pretty accurate, despite the noisiness of our neural signals. A prime mechanism for compensating for neural noise is the use of lots of neurons so that the average response can be taken, canceling out the noise.

But it may also be the case that noise has some useful functions in the nervous system. Noise could be a feature, rather than just an inconvenient bug.

In Action

To see how noise can be useful, visit Visual Perception of Stochastic Resonance (*http://neurodyn.umsl.edu/~simon/sr.html*; Java) designed by Enrico Simonotto,[1] which includes a Java applet.

A grayscale picture has noise added and the result filtered through a threshold. The process is repeated and results played like a video. Compare the picture with various levels of noise included. With a small amount of noise, you see some of the gross features of the picture—these are the parts with high light values so they always cross the threshold, whatever the noise, and produce white pixels—but the details don't show up often enough for you to make them out. With lots of noise, most of the pixels of the picture are frequently active and it's hard to make out any distinction between true parts of the picture and pixels randomly activated by noise.

But with the right amount of noise, you can clearly see what the picture is and all the details. The gross features are always there (white pixels), the fine features are there consistently enough (with time smoothing they look gray), and the pixels that are supposed to be black aren't activated enough to distract you.

How It Works

Having evolved to cope with noisy internal signals gives you a more robust system. The brain has developed to handle the odd anomalous data point, to account for random inputs thrown its way by the environment. We can make sense of the whole even if one of the parts doesn't entirely fit (you can see this in our ability to simultaneously process information **[Hack #52]**, as well). "Happy Birthday" sung down a crackly phone line is still "Happy Birthday." Compare this with your precision-designed PC; the wrong instruction at the wrong time and the whole thing crashes. The ubiquity of noise in neural processing means your brain is more of a statistical machine than a mechanistic one.

That's just a view of noise as something to be worked around, however. There's another function that noise in neural systems might be performing—it's a phenomenon from control theory called *stochastic resonance*. This says that adding noise to a signal raises the maximum possible combined signal level. Counterintuitively, this means that adding the right amount of noise to a weak signal can raise it above the threshold for detection and make it easier to detect and not less so. Figure 2-32 shows this in a graphical form. The smooth curve is the varying signal, but it never quite reaches the activation threshold. Adding noise to the signal produces the jagged line that, although it's messy, still has the same average values *and* raises it over the threshold for detection at certain points.

Just adding noise doesn't always improve things of course: you might now have a problem with your detection threshold being crossed even though there is no signal. A situation in which stochastic resonance works best is one in which you have another dimension, such as time, across which you can compare signals. Since noise changes with time, you can make use of the frequency at which the detection threshold is crossed too.

In Simonotto's applet, white pixels correspond to where the detection threshold has been crossed, and a flickering white pixel averages to gray over time. In this example, you are using time and space to constrain your judgment of whether you think a pixel has been correctly activated, and you're working in cooperation with the noise being added inside the applet, but this is exactly what your brain can do too.

End Note

1. Simonotto, E., Riani, M., Seife, C., Roberts, M., Twitty, J., & Moss, F. (1997). Visual perception of stochastic resonance. *Physical Review Letters, 78*(6), 1186–1189.

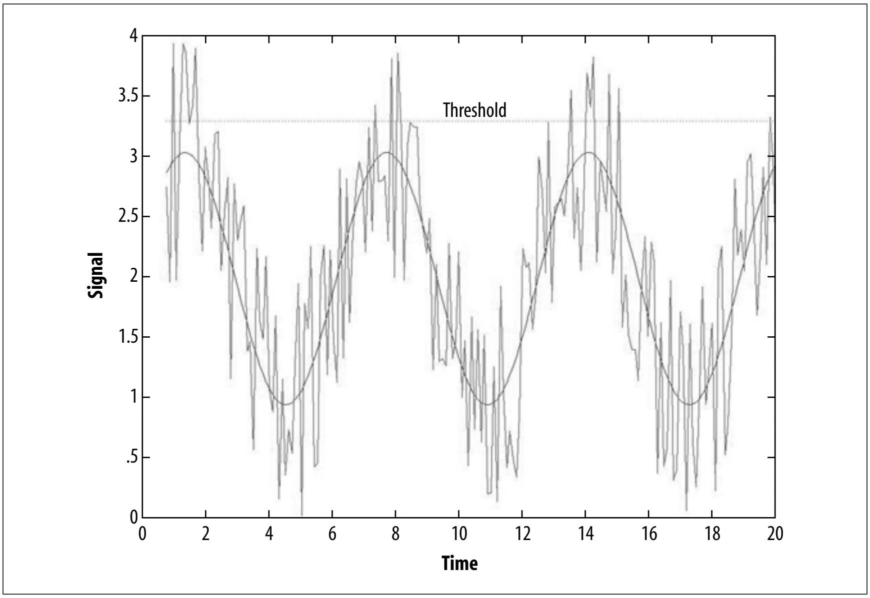

Figure 2-32. Adding noise to a signal brings it above threshold, without changing the mean value of the signal

See Also

- An example of a practical application of stochastic resonance theory, in the form of a hearing aid: Morse, R. P., & Evans, E. F. (1996). Enhancement of vowel coding for cochlear implants by addition of noise. *Nature Medicine, 2*(8), 928–932.