Temporal Scaling of Neural Dynamics in the Motor Cortex of the Singing Mouse: The Backstory

By Arkarup Banerjee

As scientists, we usually publish a fully worked-out idea and shy away from discussing the *process* of doing science. What appears to be a straightforward progression from a clear question to a satisfying answer is usually a circuitous path marked by false starts, wrong turns, and dead ends. How science is *actually* done often gets lost in the narrative simplicity of a published paper. As a case in point, the official narrative of this study could not be more different from how it actually happened. For instance, the computational model, which ended up as the last figure of the manuscript, was conceived before any of the data was collected. This is my attempt to share the backstory.

Why would you be interested? I'm not sure. Perhaps you're a science enthusiast or curious reader who struggled through the complexity of the scientific paper but, after reading this, finally understands what the fuss is all about. Or you're a student, and reading this gives you the confidence that you can do it too. Or maybe you're simply intrigued by the story and you find this mildly entertaining. If any of those apply, then this attempt has succeeded.

Singing Mouse
in an operatic posture

The story of this study ostensibly begins in March 2019. I was attending the annual meeting of Computational and Systems Neuroscience in Lisbon. My first paper (with Daniel Okobi, Steve Phelps, and Michael Long) on the neuroscience of the singing mouse had just been published as a cover article in the Science magazine. What is a "singing" mouse, you ask? These are exotic rodents from Costa Rican cloud forests that are remarkably good at vocalizing in a "call-and-response" fashion with each other, similar to how we "take turns" in a conversation. They adopt an operatic posture (up on their hind legs, snout pointed towards the sky) and belt out human-audible songs that last many seconds, following a stereotyped sequence of hundreds of notes, reaching a loud crescendo before trailing off. Simply put, it's a mouse that sings like a bird!

Steve Phelps at U.T. Austin pioneered work on these singing mice, focusing on the genetics, ecology, and neuroscience of vocal communication for about 15 years. A few years before I joined, my postdoctoral advisor, Michael Long, established a collaboration and brought singing mice to NYU. Michael's lab had primarily worked on songbirds (like zebra finches) that produce melodies when wooing a female. Over the previous 40 years, a lot had been discovered about the motor control pathways in zebra finches. Michael wanted to investigate whether the neural algorithms identified in a bird brain would generalize to a mouse with a six-layered neocortex, just like ours. Particularly enamored with this question, I joined the lab to study vocal behaviors in this intriguing system.

Collaborating with Daniel Okobi, a graduate student, we identified a region of the motor cortex crucial for this behavior. Without it, singing mice could still vocalize but were unable to flexibly adjust their songs to suit conversational needs. This finding was a striking demonstration of a motor cortical region controlling vocal behaviors in a rodent—previously thought to be the exclusive domain of primates. Our article received considerable attention in the media and the scientific community. It was thrilling, but it also raised many more questions than it answered. Foremost among them: How do neurons in the motor cortex control vocalizations in singing mice? We had no idea how this neural mechanism worked.

On my flight back from Lisbon, I decided to tackle this question head-on. First, we had to solve a technical challenge: how to record neural activity in these animals while they were singing. This was partly an engineering problem, since singing mice are tiny (around 15 g). We worked with a company called Diagnostic Biochips to design a lightweight 128-channel electrical probe, made from silicon wafers, that could be implanted into the mouse brain. There was also the challenge of mastering the surgical procedure of implanting these electrodes. It took many months of trying and failing before I got good at it. After persistent nudging from Michael and help from a graduate student, Margot Elmaleh, I finally managed to record spikes from motor cortical neurons using an implanted array of silicon probes. Seeing those “first spikes” was both exhilarating and relieving—the first ever recorded in this species!

While I was honing my experimental skills, I felt intellectually stagnant and decided to build a “toy model” mostly for my own amusement. Although I’m trained as an experimentalist, I’ve had a knack for modeling for as long as I can remember, a habit picked up by working with evolutionary biologists and physicists in my formative years.

Quantitative modeling is challenging, fun, and almost an art form. The first step is to strip away unnecessary complexities and focus on the core idea—a “spherical cow” approach (if you don’t know the phrase, look it up—it’s worth it). After pruning the details, you create a simple-enough mathematical model (usually with equations) to explain it. That may sound easy, but nobody can agree on which aspects are “details” and which are the essential phenomena you should focus on. A detail a theorist finds trivial can be hard-earned data for an experimentalist, leading to endless, if usually polite, debates. In this project, though, it was less contentious because I was both doing the experiments and building the model!

To appreciate the knowledge gap we wanted to address, here’s a bit more about the motor cortex and its role in vocal production. We knew singing mice take part in rapid vocal interactions, where the responding animal can significantly change the song’s tempo and duration—sometimes drastically. For instance, in under a minute, a mouse might sing a 7-second song, then follow up with a 4-second tune, then finish with a 16-second epic. We also knew that silencing the motor cortex with a drug quenched this behavioral range; drugged animals could only produce songs of similar durations (around 6-7 seconds), but not much shorter or longer, making for a monotonous “conversation.” Furthermore, physically lowering the temperature (cooling) of the motor cortex was enough to stretch song durations—colder cortex, longer songs, faster tempos. These experiments demonstrated that neural activity in the motor cortex controls song duration.

So, there it was: our core phenomenon. How does the brain, particularly the motor cortex, generate such flexible behavior? A satisfactory model had to account for how neurons change their firing rates to control both song duration and tempo, while aligning with the evidence from cortical silencing and cooling. I struggled with this problem for a few days, thinking about it constantly—during office hours, on my commute from the suburbs of Long Island to New York City, even at home before bed. I had sketches on napkins in our kitchen. Luckily, my wife is also a scientist and is pretty tolerant of these quirks. After a couple of days, I realized that a slight twist on a classic neuroscience concept—the "integrate-and-fire" model—was exactly what I needed to make it work.

The integrate-and-fire model was introduced in 1905 by the French scientist Louise Lapicque to explain how neurons generate action potentials (spikes). A neuron's membrane voltage sits at a resting level, then "integrates" inputs from other neurons, causing its voltage to climb until it reaches a threshold. At that point, the neuron fires an action potential. Afterward, the voltage resets to the resting value, and the cycle repeats. Amazingly, this model was conceived long before the biological intricacies of multiple ion channels were discovered, yet it remains deeply relevant in neuroscience.

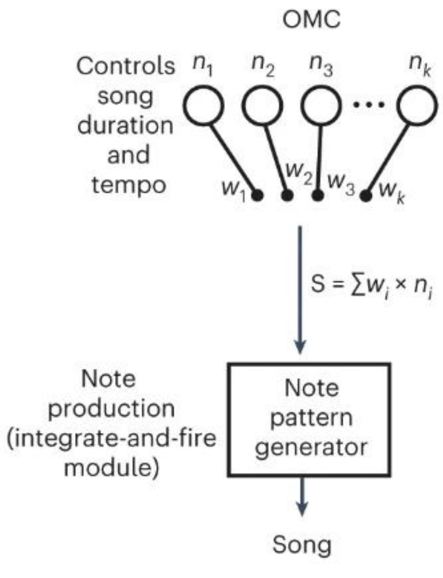

Conceptual model of the hierarchical motor control

So how does this help explain vocal flexibility in singing mice? Picture a two-step motor hierarchy that generates the song. At the lower level is a motor command neuron; each time it fires, the mouse produces a single note. The higher-level motor cortex sends inputs to this command neuron. These inputs are "integrated" (just as the integrate-and-fire model predicts) to generate the action potentials of the command neuron in a temporal sequence—thus creating a song. Here's the key insight: if the dynamics in the motor cortex simply stretch in time by a constant factor (let's call it "c"), the song's tempo also scales by that same factor "c", and the duration automatically adjusts. Thus, through this formulation, the speed of neural dynamics in the motor cortex could directly control song duration and tempo. The model predicts that motor cortex neurons just need to "stretch" or "compress" their firing rates over time, and the degree of that scaling determines each song's characteristics. This linear scaling of neural activity over time (temporal scaling) allows for a broad range of motor flexibility by simply modulating the speed of the cortex's internal dynamics. I realized this suddenly one afternoon in the lab and rushed to find my colleague and friend Robert Egger, a physicist by training. I explained the model, and he agreed that the argument seemed solid and should work. We had previously worked on a paper to highlight how mild cooling can be used in neuroscience to discover neural algorithms, so he was able to quickly verify that this model would immediately explain the cooling data in the singing mouse. I was excited because I'd made a prediction about what would happen in the brain before I actually recorded any neuronal data!

The rest was conceptually straightforward but technically laborious. Eventually, I managed to implant high-density silicon probes into the brains of singing mice and recorded neural dynamics during songs of various tempos and durations. When I plotted an example neuron's firing rates across 30 different songs—sorting them from the shortest (top row) to the longest (bottom row)—the temporal scaling popped out of the plot. The firing patterns for long songs at the bottom looked exactly like the shorter ones at the top, only stretched like an elastic band. This was the telltale signature of temporal scaling that the model predicted. I immediately knew we were onto something.

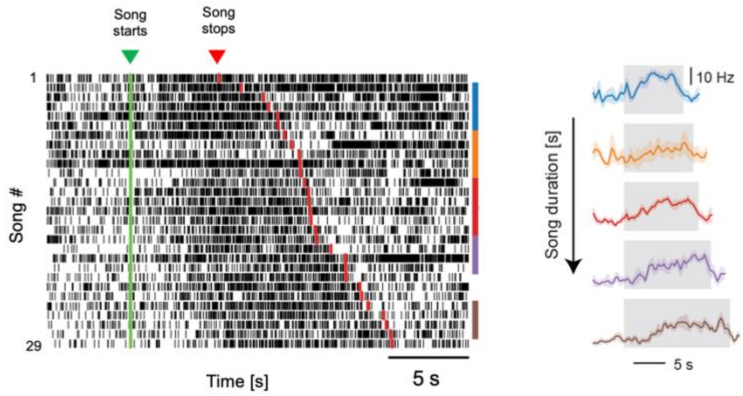

Spiking activity (raster plot, left) of the very first neuron I found to show time-scaling. Each black tick is an action potential. Notice how the firing profile (right) of longer songs looks more like a stretched version of

I showed the figure to Michael, who took a measured approach and told me to run more rigorous analyses after collecting additional data. In collaboration with our fantastic colleagues Feng Chen and Shaul Druckman at Stanford University, we conducted extensive analyses on the data and confirmed that individual neurons in the motor cortex indeed varied their "speed" on single trials to dictate the song tempo and duration. Looking back at the literature, we realized that this temporal scaling algorithm has been previously demonstrated in cortical activity of macaques trained to judge the passage of time (check out work from Jazayeri lab at MIT). We suspect that this mechanism—adjusting the speed of cortical dynamics to produce motor flexibility—may apply broadly to hierarchical motor systems across many species. This could represent a canonical computation in the brain.

There is a lot more in the paper – we discovered that single-trial neural variability is affected by the singing behavior, that there are approximately eight types of neural activity profiles, all of which scale in time, and that motor cortex dynamics also reflect faster time-scale processes potentially related to sensory feedback. From the initial flutter of an idea to a finished project, it took us four years! It was really a pleasure to work on this study with my coauthors. Good science usually takes time, way longer than you expect. I heard it from a senior colleague that a good rule of thumb is to multiply your best prediction of how long something should take by a factor of three and then go up by an order of magnitude (i.e., a week become 3 months).

This study has raised even more questions: What determines the speed of cortical dynamics? What is the cellular and circuit mechanism behind the integrate-and-fire algorithm that the model proposes? My lab is currently investigating these questions, and we’re beginning to form ideas about how they might work. Still, as I just mentioned, it will likely take years before we really know the answers.