a firm yes. The agent then glanced at the SUV and at us sitting in the car. He then backed away from the car and waved us to go ahead. The SUV began moving past the checkpoint. We drove slowly, and some of us turned around and watched our second SUV stopped at the checkpoint. Eventually, the second SUV was also allowed through. I was relieved.

We broke out into conversation. For about half of us, this was our first time going through Border Patrol. I was so happy that it was so easy. But at the same time, I was worried about how it could've been different. I thought about our privilege. I thought about how we could easily pass from the contested border back to the US because of who we were. We were the right type of citizen scientists. We were mostly white. We drove a newer SUV bearing the symbol of our prestigious university. We did environmental science work—for free. We were the right type of citizens, and this was determined by easy passage by the Border Patrol. Maybe we should re-examine what we mean by *citizen* in citizen scientist. Does it mean safe passage through nation-states? Does it include undocumented immigrants? Does it include historically and currently marginalized groups that aren't allowed to be citizens? Maybe we should call citizen science something else.

The conversation then grew more divisive. Some members grew up with Border Patrol and had family members in their ranks. Some members were first-generation, whose parents came from Mexico or South America. Ideas and ideologies clashed. I, too, was first-generation. Yet, the conversation surrounding Border Patrols and walls was unfairly focused on a small group of immigrants. I pointed out how bizarre the idea of a border was. How strange that a river could be so contested. I thought back to when we were on the Rio Grande, and how you wouldn't necessarily know that you were on US land, and that the land across the river was Mexico. How some days you could get turned around and lose track; because on the river it didn't really matter which side was the US or Mexico. But now, we have people who have never been to the border enacting policies presenting certain citizens from crossing from one side to the other.

We came to the Rio Grande as citizen scientists to monitor water quality. We took a risk. We all made it back. But many of us, including me at the time, didn't realize what it meant to do citizen science work on the border. I didn't know that we'd be faced with privilege. The privilege of getting to do scientific work. The privilege of getting *to come back*.

The People and Serendipity of the EyesOnALZ project

Pietro Michelucci

Timing is everything. It was during the pre-Web, Internet period of the late 80s that I was using so-called "connectionist models" (now known as "artificial neural networks") to see how well they could explain human behavior. At the time, I was struck by the parallel between neurons in the brain and people connected to the Internet. To me, they both seemed like examples of computational nodes joined by a network, with the potential to produce a collective output greater than the sum of its parts. I could only wonder about the potential capabilities of systems that somehow combined the thinking of thousands or millions of networked humans.

At that time, my academic mentor espoused a contrarian view about knowledge representation in the brain: that two very different kinds of information processing are necessary to account for human behavior. I eventually realized that one of these types—symbolic reasoning—seemed better suited to machines than humans. This suggested the value of creating partnerships between humans and machines.

I might have predicted those formative experiences would lead to a life-long pursuit of human computation—an emerging field that leverages the complementary abilities of networked humans and machines to solve real-world problems. There is no way I would have guessed the first such problem

I tackled using human computation would be Alzheimer's disease, which turned out to be coincidentally appropriate, as it employs human minds en masse to save human minds en masse.

Janis Dickinson, a professor of natural resources, happened to meet Chris Schaffer at a Cornell University faculty dinner. Schaffer, a professor of biomedical engineering, described his Alzheimer's disease research to Janis who, realizing there might be an opportunity to apply human computation, made an introduction. The mere prospect of addressing a disease like Alzheimer's, which has no effective treatment or cure, compelled me to follow up immediately. What happened next can only be described as an alignment of the stars.

I met with Chris at his Cornell office. He offered me a laboratory-grade espresso and then we sat down as he began to describe his research. They were studying Alzheimer's by inserting the human gene for the disease into mice and then comparing outcomes to the "wild type" mice without the disease. He went on to describe a new imaging technique they had invented that allows them to see blood flow in the brains of these mice. This enabled a key finding—that the Alzheimer's mice exhibited a relatively high rate of stalled brain capillaries, meaning that many of the tiny vessels were plugged and had no blood flowing through them.

Chris and his colleagues conducted a follow-up study showing that, due to downstream effects, these capillary stalls were responsible for an overall reduction of 30% of blood flow in the mouse brains, which is the same reduction observed in humans with the disease. Chris explained that although we've long known that human Alzheimer's patients have reduced brain blood flow, we've never understood why. The new imagining technique created the opportunity to investigate this phenomenon in Alzheimer's mice exhibiting the reduced blood flow.

This enabled a serendipitous discovery. Victorine "Torie" Muse joined the Shaffer-Nishimura Lab as an undergraduate research assistant. In one of her first assignments, she was instructed to use an antibody as a marker to help identify stalled blood vessels in the mice being studied. At first, Torie was discouraged when she didn't find any stalled vessels, thinking that perhaps she had made a methodological error. Despite feeling embarrassed, Torie brought her negative findings to Chris, who realized the antibody intended to help visualize the stalled vessels was actually interfering with the stalling mechanism, suggesting an immunological aspect to the reduced blood flow. Thus, Torie's accidental finding led to the first understanding ever of the mechanisms underlying reduced brain blood flow in Alzheimer's disease.

Based on this key finding, additional mouse studies showed that unclogging these capillaries restored brain blood flow to normal levels and reversed cognitive symptoms such as memory loss. The problem was the drugs used to restore the blood flow also compromised the immune system, which would eventually kill the patient. In other words, the cure was worse than the disease. This left Chris and his laboratory with the goal of finding a new drug that would restore blood flow safely and effectively.

I couldn't believe my ears. It sounded like Chris was telling me they were on the verge of finding the first effective treatment for Alzheimer's disease. But it turns out there was a catch. Analyzing the mouse brain images to find the stalled capillaries is extremely challenging. He said for each week's worth of data collection it would take a trained laboratory technician six months to a year to analyze the data. He went on to say even though their findings were promising, because of this analytic bottleneck, it could take decades to analyze the data needed to arrive at a treatment target.

My colleagues and I believe in a moral imperative to not waste human time solving a problem that machines could address more effectively. So my next question was, "have you tried using machines to solve this problem?" It turned out the laboratory had already pursued the best available machine-based methods, which turned out to be grossly inadequate for achieving the needed data quality. I asked Chris to show me how the lab technicians performed the manual data analysis. As soon as I saw the videos of the blood vessels in the mouse brains I instantly felt déjà vu. I was reminded of a citizen science project I had participated in a decade earlier called stardust@home.

In the stardust@home project, space scientists were trying to detect interstellar dust particles embedded in aerogel that was brought back to earth by a satellite flown through the tail of a comment. This required looking through a million microscopic images of aerogel, which the scientists estimated would take approximately 100 years. They created an online activity that allows volunteer participants to look through a virtual microscope to try to find the Stardust. With over 30,000 so-called "dusters" participating, they were able to sort through one million images and find seven interstellar dust particles in just a few years. A finding reported in the journal *Science*.

Recalling my stardust@home experience, I wondered if the virtual microscope could be adapted to look at blood flow in mouse brains. I got so excited about this that I phoned the project leader, Andrew Westphal, and told him about the Alzheimer's research and how we might be able to use his virtual microscope interface to help speed up Cornell's research. At first, he was skeptical but when I showed him the images he agreed that it could work. Having lost his father to Alzheimer's disease, Andrew was also personally motivated and graciously joined our team, offering to adapt his platform to the Alzheimer's data.

Soon thereafter, Rod Corriveau, a program officer at the National Institute of Neurological Disorders, heard me describe these ideas in a brief talk at NIH's National Alzheimer's Summit. Recognizing the unconventional approach, he subsequently introduced me to Guy Eakin, who was then the Scientific Director at BrightFocus Foundation, an organization with a high success rate for giving wings to highly innovative research. Guy was aware of citizen science and appreciated its research potential, so he invited a proposal. BrightFocus awarded two consecutive grants to develop what would become known as the EyesOnALZ project and the Stall Catchers citizen science game. In helping us test an early prototype, Guy eventually gained the distinction of becoming the first "catcher" to perform an in-flight annotation.

The dream team of collaborators included the stardust@home research contingent and cross-over volunteers (our number one ranked catcher today, Mike Capraro, is a former duster). The Cornell contingent, run by Chris Schaffer and Nozomi Nishimura, included their PhD student, Mohammad Haft Javaherian, who invented machine learning methods that enable the the drawing of outlines around target blood vessels in Stall Catchers, Oliver Bracko, a post-doc who conducted many of the Alzheimer's studies, and Torie Muse, the undisputed blood vessel expert, who trained other lab technicians to analyze the blood vessel images and established "ground truth" answers we could use to validate the platform. We also received great support from veteran citizen science practitioners from the EyeWire project.

Due to the heroic efforts of many people, Stall Catchers officially launched on schedule on October 1, 2016. Stall Catchers was hacked four days later, when three users legitimately climbed the leaderboard, and then, once in the top positions, changed their usernames to create a prominent political message in all capital letters that all the other users would see. This relatively benign experience turned out to be a useful lesson by assuming that "if it can be hacked it will be hacked," and thus taking preventive measures. Despite our best efforts, Stall Catchers was not, however, as bulletproof as we thought.

A year and a half later we were contacted by Firas Khatib, one of the developers of fold.it, who was teaching a university course about citizen science and using Stall Catchers as an example. Firas contacted me to let me know that some of his students were skillful hackers and that even though he warned them not to hack Stall Catchers, we should probably remain vigilant anyway. This seemed like a great opportunity to test our defenses, so I suggested that he invite his class to hack our system, and if any students were successful and shared their methods, we would acknowledge them in a relevant publication. Firas went one step further and offered them extra credit if they could do it. As it turns out, Noato Nicol, one of the students, succeeded in hacking Stall Catchers and shared his methods, allowing us to protect the system from future related attacks.

As we designed new features for Stall Catchers, the notion of "Gamification"—taking a mundane task and introducing game-like elements to make it more entertaining—was central to our approach. We joined forces with a software project manager from the gaming industry, who helped us keep up with a swelling platform and a growing appetite for new features. We also recognized the need to develop a strong community. We were tickled to be joined by Alice Sheppard, who was instrumental in developing GalazyZoo's community forum, one of the largest in history with over 650,000 posts. Alice recognized that a key success ingredient for Galaxy Zoo was to provide growth opportunities for community members. Based on this wisdom, she invited several of our most active community members to become mentors who help others and alert us to any issues. One particularly active mentor is Guy Calkins, who ranks second on the overall leaderboard at the time of this writing (exceeded only by Mike Capraro). Guy has not only alerted us to numerous data issues, but keeps close tabs on the forum, answering when he can, and passing questions along to the project team when he isn't sure. He has also been an honored guest at our headquarters, where he proposed carefully conceived process improvements and feature ideas. Indeed, with his support for a live chat box, we felt confident moving forward with this new community feature that will be released in Stall Catchers by the time this article is published.

Once we launched Stall Catchers, we ran a validation study to ensure we could extract expert-like answers from the crowd. In a single month, almost 1000 volunteers analyzed 96,000 vessels in the game. The results exceeded Cornell's research requirements with crowd answers achieving over 95% sensitivity and specificity, the levels needed to draw strong conclusions from the data.

Next, we focused on growing the participant base and making more efficient use of each annotation in the game. By the end of 2017, volunteer recruitment was fueled by online outreach, public booths at prominent local events such as the USA Science & Engineering Festival, the British Science Festival, and GameOn, a featured segment in a PBS documentary mini-series called "The Crowd & The Cloud", and our first international catchathon—a synchronized online event with 20 teams from 15 countries and six continents, including sub-Saharan teams like "Ugacatchers" who competed from Uganda with low bandwidth conditions. How do you get 20 teams from around the world to participate? Easy, you buy them pizza—a lesson we learned from Erin Lamichhane, a middle school teacher who used that incentive to entice her students to come into school on a Saturday to clinch their victory in our month-long team competition.

During that time, the Stall Catchers community grew to 8,000 users. We also developed improved methods to extract the "wisdom of the crowds". In Stall Catchers we collect many answers about the same blood vessel and then combine them to produce a single expert-like crowd answer. This approach allows any individual errors to get washed out, so volunteers can rest assured that the data quality will be high even when they aren't at the top of their game. These improvements allowed us to reduce the number of individual answers we had to collect for each movie from 20 to approximately seven. The combination of more users and efficiency improvements increased our analytic throughput so that it was consistently double that of the lab, effectively cutting the research time in half. Our goal, however, is to reach a ten-fold improvement over the lab in order to reduce the potential time to a treatment target down to just a few years.

Not all of our players have been able to stick around. One of our most dedicated catchers during the first year was a woman in her late 80s who would play long enough each day to remain in the top 20 on the leaderboard. One day we noticed she had not been active, and then we suddenly received an email in which she graciously alerted us of a personal injury that would interfere with playing the game. We could hardly believe that under the circumstances she had gone out of her way to explain why she was absent from Stall Catchers. Our main concern was her health, and we sent her a get well card with a heartfelt message of gratitude from the team. This sense of mutual dedication to a shared enterprise in service of a common goal pervades our community and instills in our team a great sense of loyalty toward all of our catchers.

Erin Lamichhane, a technology teacher from Idaho, organized the winning team of 250 students for our #CrushALZ team challenge that accomplished 8 months worth of lab research in just 4 short weeks. Erin's students created their team logo and self-organized to recruit more players doubling the size of their team over the competition period. I visited Erin and her students in 2017 to award her middle school the world's first citizen science trophy, marking a new era for citizen science in the classroom. Before leaving, I had the privilege of spending a few hours with the students, who sent back a wish list of features that have since been integrated into the prioritized schedule that drives our platform development. They were so excited to imagine their ideas showing up in the Stall Catchers game, and we were equally excited to add their great ideas to our list—ideas these young minds generated that will help cure Alzheimer's disease.

When Citizens Do Science

Jaden J. A. Hastings

It was a serious illness that forced me to leave behind an emerging career as an academic scientist. I became a shut-in; unable to spend time away from the house, yet I needed an outlet for my intense curiosity and creativity. This was the original impetus for setting up a laboratory in the laundry room next to my bedroom in 2009. Soon after, I discovered the growing international biohacking community—present on email lists, meetups, and online forums—and began connecting with colleagues as well as aiding in the formation of new community laboratories for other biohackers.

I prefer to use the term "biohacking" or "independent research," rather than "citizen science," as they are conducted in different ways. "Citizen science" is typically research initiated, coordinated, and funded by a researcher, or network of researchers, based in an institution and seeking the assistance of the public for data gathering or distributed analysis. Biohacking, or independent research, is initiated by a private individual who relies on their own resources and initiative to conduct their scientific enquiry. In my story, biohacking refers to those activities where the layperson employs a mixture of lo-fi and hi-tech tools within the confines of a private residence or shared hackspace to pursue research of their own volition and interest within the scope of biotechnology.

In my mind, to embrace biohacking is to accept this first principle: that we all possess an intrinsic drive for discovery and exploration. Certainly, there are a few who wish to devote themselves professionally to the pursuit of a narrow course of inquiry in great depth and are fortunate enough to find academic positions that can provide adequate support for their research. There remain precious few such funded positions available, and yet so many of us remain voraciously curious and seek a means through which we can continue to contribute toward the long arc of emerging scientific knowledge.

Contemporary biohacking possesses deep roots in the history of science, particularly those practices that pre-date formalized research within the walls of institutions and government-funded research. The concept of biohacking—of utilizing the tools of biotechnology outside of the purview of commerce or academia—emerged once individuals with sufficient knowledge of the techniques were able to procure the requisite equipment and reagents from myriad sources. Building upon the competencies and praxes of their infosec hacker predecessors—indeed, with some overlap within their ranks—the sophistication and impact of biohacking has developed in concordance with the (a) decreased price and increased availability of common reagents and equipment, along with the (b) rapid and open dissemination of scientific knowledge.

My own home laboratories—I have set up more than one as well as aided in the start-up of more