Sex, Flies, and Videotape

By Douglas L. Smith

A lunging fly looks angry, but the "emotion" is a hard-wired reflex. Now a computer-vision system that watches flies and can figure out what they're doing is helping biologists trace those wires—and might one day be able to read human emotions.

How deep does emotion go? When fear trumps reason, pop psychologists tell us it's the older, reptilian part of our brain taking over. But what brain circuits are really responsible, and how are they wired up? To answer this very big question, a Caltech neurobiologist, a computer scientist, and a bioengineer have started very small—with *Drosophila melanogaster*, otherwise known as the common fruit fly.

The fly's tiny brain contains only about 40,000 nerve cells—not counting the optic lobes, which include another 60,000 or so—and, to a first approximation, is more or less hardwired, says biologist David Anderson. "We do not have anything *close* to a complete wiring diagram, but when you look at the major branches in a neuron's dendritic tree and where those branches go, they seem to be remarkably constant from one fly to another." So, like a person in a strange house flipping the switches in the front hall to find out where the lights are, Anderson is turning specific nerve cells on and off to see what happens.

The seeds of Anderson's work were planted at Caltech in the 1960s. While researchers elsewhere were teaching mice to run mazes, biologist Seymour Benzer began mutating flies to induce behavioral oddities—establishing what has become a very fertile field. As flies have a generation time of only 12 days (versus some nine weeks for mice), the experiments proceeded at a gratifying pace, and his lab discovered a host of genes responsible for controlling such things as how flies responded to light, when they slept, and whether they mated. Although Benzer formally retired in 1992, he continued working right up to his death in 2007.

When Anderson got interested in the neural-circuitry problem about a decade ago, he intended to study mice, and he and biologist Henry Lester began developing a set of techniques to turn mouse neurons on and off. "But I had always had fly envy," Anderson says. "I spent a lot of time talking with Seymour about whether one could study primitive versions of emotion-like behaviors, like fear, in flies. And there are a lot of fancy genetic manipulations that you can do in flies that you can't do so easily in mice." Around 2005, "after spending three or four years really struggling to get the mouse system under way, we decided to initiate a program with flies. So the reason I started with them—in addition to sheer impatience—is the potential to screen thousands of lines of genetically altered flies to look for behavioral changes."

FEAR AND AGGRESSION IN PASADENA

Anderson also eventually switched his focus from fear to aggression. "When people see videos of flies fighting, they don't have to be convinced that this is analogous to an emotional behavior in humans. It's harder to show that a fly is afraid of something. There's a quote from Charles Darwin's 1872 book, *The Expression of the Emotions in Man and Animals*, that I love: 'Even insects express anger, terror, jealousy, and love by their stridulation.' Now, this doesn't mean that I think that flies get angry when they fight, or feel anything like anger. When we talk about 'emotional behavior' in animals, we're talking about an observable motor behavior, not a subjective state that might accompany such behavior."

These subjective states, however, are very much on the mind of machine-vision expert Pietro Perona. Since the mid-1990s, Perona has been trying to figure out how to teach computers to look at people, read their emotions, and divine their intentions—something humans do instinctively. (See "The Machine Stares Back," *E&S* 1999, Nos. 1/2.) Besides leading to some truly awesome video games, such a system could be used by market researchers to find out what people really think of a product being pitched to them, or by security cameras to decide whether that nervous-looking fellow in the ATM queue is a potential armed robber.

But when an action as simple as walking looks so different in front and side views, good luck trying to codify the distinction between a joyous strut and a furtive skulk. "How can I decompose human behavior into meaningful motions, which are its fundamental ingredients?" says Perona. "I realized at some point that I had to study behavior in a simpler setting, and David convinced me that flies have an enormous repertoire of interesting behaviors. Flies are very simple animals and their bodies are extremely easy to track. You can put a camera on them and watch them as they fight, as they court, as they mate, as they look for food."

Anderson was also talking with Michael Dickinson, a zoologist-turned-engineer who had been deconstructing *Drosophila*'s flight-control systems with an eye toward designing free-flying microrobots. (See "Come Fly with Me," *E&S* 2003, No. 3.) A fly can flap its wings 250 times per second, and it can change its course by 90 degrees in 50 milliseconds, so Dickinson's lab was shooting high-speed videos of individual flies on the wing. A computer processed these videos to reconstruct each wing's motions in three dimensions, but only after a live human had gone through the footage, one frame at a time, to trace the outlines of the wings' silhouettes. But Dickinson's interests were shifting from the fluid dynamics of flies in midair to the group dynamics of flies in large numbers, which would mean tracing dozens of flies at once.

Both Anderson and Dickinson envisioned the same strategy: put some flies in an "arena"—a rather grandiose term for an enclosure that might range in size from a postage stamp to a dinner plate—mount a small camcorder overhead, like the Goodyear blimp looking down on the 50-yard line at the Rose Bowl, and ask a computer to report on what the flies are doing. The computer, working tirelessly, would analyze thousands and thousands of hours of video to create vast databases of fly behavior that could be mined statistically.

But the two sets of software specs that resulted were nearly polar opposites. Anderson needed to watch pairs of flies in close quarters and catalog each occurrence of any of an assortment of predefined aggressive or courting actions—Perona's "meaningful motions," which, when performed in various sequences, add up to the complex behaviors we put names to. Dickinson wanted to follow 50 individual flies out in the open simultaneously, keeping each fly's identity straight while mapping its path. His software needed to classify the fly's movements into a few broad categories—Was it walking forward? Backing up?—but the main purpose was to reveal patterns of social interaction.

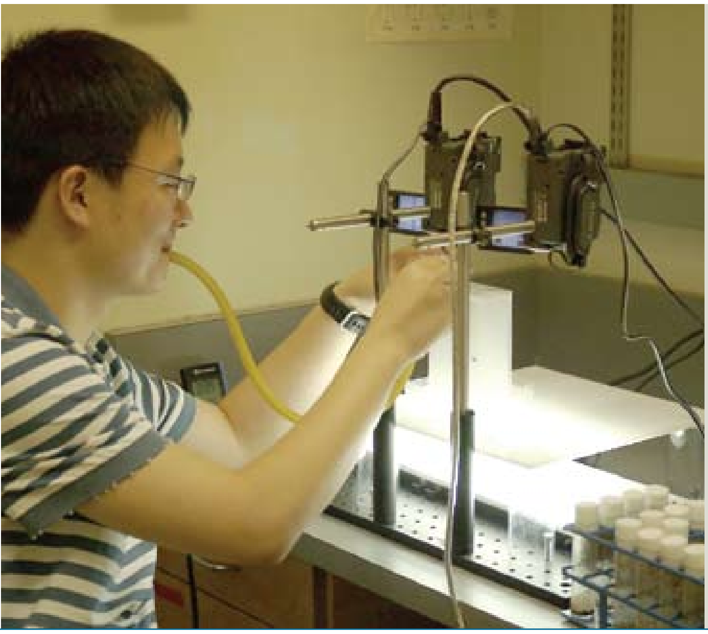

Anderson's grad student Liming Wang stocks an arena—the tall rectangular container underneath the leftmost camcorder—with a pair of flies from one of the vials at right. A fluorescent lightbox illuminates the arena floor from below, creating a clean, uniform background that makes it easier for the computer to find the flies. He is ferrying the flies with an aspirator, which is basically a length of surgical tubing with a plastic tip.

EYES ON THE FLIES

Back in 2005, the state of the art in behavior tracking owed more to the sweatshop than the supercomputer—grad students and postdocs watching endless hours of grainy video and making tick marks on pieces of paper whenever fly X lunged at fly Y. Each video was only 20 minutes long, but if you factored in all the scrolling back and forth to mark the exact frame where every lunge began and ended, then completely cataloging just the lunges could take well over an hour. If you also wanted to count the chases, or the touches, you had to go back and rewatch the entire video all over again.

Not only is this mind-numbingly boring, it's error-prone. Besides the obvious ways

WHAT IS "BEHAVIOR," ANYWAY?

Perona's ultimate goal is to build a computer that interacts with us the way we do. "We look at each other, and I know you're sitting comfortably and I know you're paying attention to me, taking notes," he explains. "That's very useful for me and I would like a machine to be able to do the same." Perona picks up all this information intuitively, with a glance at body language and facial expression, but a computer needs things to be spelled out for it. Figuring out what we need to tell the computer brings up a very basic question—what do we mean by "behavior" in the first place?

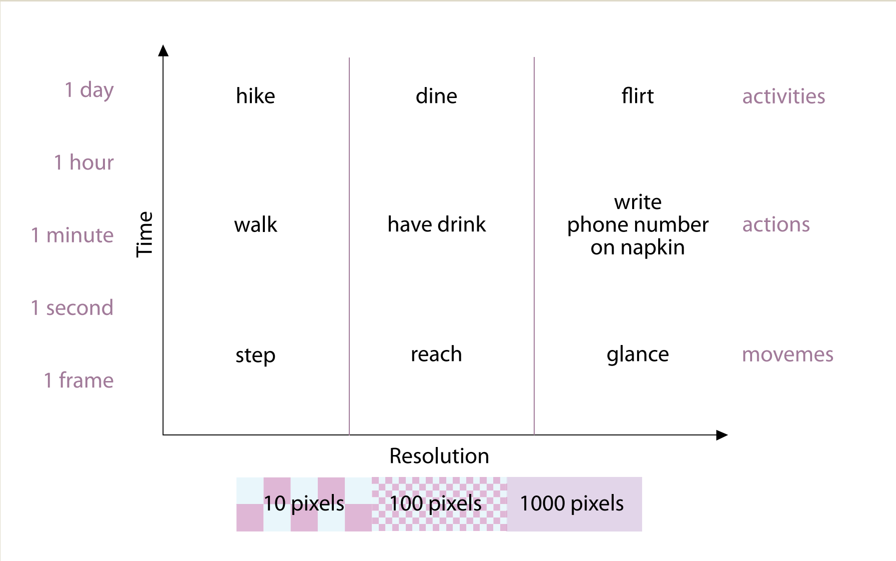

It turns out that we mean many things, as the chart at right shows. We tend to think in terms of long-term behavior, which we describe at a very high level. "That man is out to pick a fight," a pool-hall bouncer might say. But this abstraction is based on observing concrete activity—the belligerent fellow has been getting drunk for the past hour, and is now sliding off his bar stool and heading for the biggest dude in the room. Activity, in turn, is a collection of individual actions that might take a minute or less each—chugging a beer, for example. And at the very lowest level, an action can be broken down into a sequence of movements—reaching for the mug, gripping the handle, rasing the brim to the lips, tilting back the head and arm, and slamming the empty down on the bar.

These are the "meaningful motions" Perona is looking for, which he calls "movemes" by analogy to the phonemes that are the fundamental sounds of speech. A lisped "s" is still an "s," and someone speaking English with a German accent may inject a little extra phlegm into the "ch" sound, but the meaning is still clear. Similarly, a reach in any direction is still a reach, even if it's a sloppily executed one that winds up spilling half the mug.

In other words, movemes are the building blocks of behavior that can be described in such a way that they make sense to a computer, when all the machine has to go on is a video feed that it can examine frame by frame to see if any of the pixels have changed.

This brings us to the chart's horizontal axis—seeing a behavior takes various amounts of pixels, as well as different lengths of time. Some movemes will be obvious even at very low resolutions. "If you punch somebody," says Perona, "I may see that even if your whole body only appears in 10 pixels. You could be very far away, just a ghost in the distance. But if you wink at me from across the room, I may need to put a thousand pixels on your body to be able to see that eye motion. So there are multiple scales of resolution in time and space that are meaningful to us."

Substituting fruit flies for humans has allowed Perona to explore this notion of movemes using a creature with a much smaller vocabulary of gestures. "We have made quite a bit of progress in understanding how to think about these problems," he says. "What is the signal that is there in images? How do you harvest it? And how do you decide if something is happening?"

When these problems are eventually solved, Perona says, we could wire up "smart homes" for the elderly who live alone, where a computer automatically calls 911 if you've fallen and you can't get up, or summons your doctor if you look a little green. Factory floors and construction sites could be made safer, since "every move of every worker could be followed and evaluated for risk, and the worker could be briefed at the end of the day on better safety practices." And, of course, there'd be the killer iPhone apps. "Your cell phone might one day be able to tell you what's wrong with your golf swing, and maybe give you tips on your tennis backhand as well." *—DS* E&S

Right: Top and side views of three things CADABRA is programmed to recognize. Going from left to right, we see a lunge, a tussle, and a wing threat.
Below right: The computer extracts the fly's image from the background pixels, then fits an oval around the fly's body.
Bottom right: Some of the parameters CADABRA calculates for every encounter between a pair of flies. The scale bars are one millimeter long.

All figures from Dankert, et al., *Nature Methods* 6: 4 (2009)

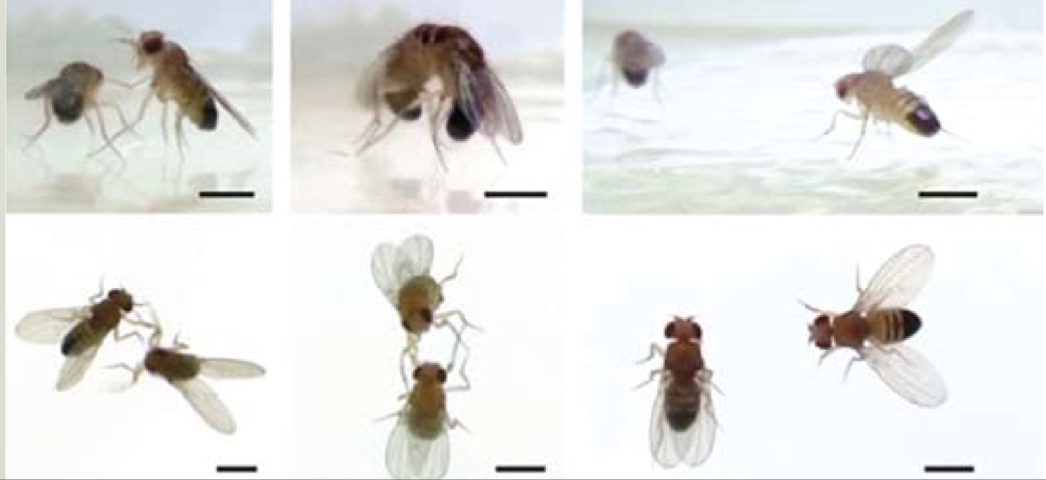

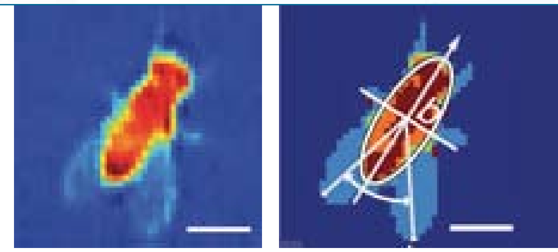

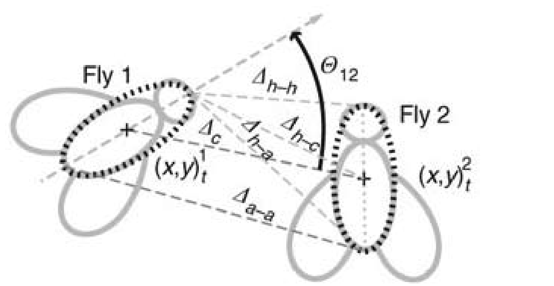

to miss something—an ill-timed sneeze, a ringing phone, or simply zoning out for a few minutes—behavior can be in the eye of the beholder. Says Perona, "How long does one fly have to follow the other for it to count as 'following'? At what distance should they be? How fast should they be going? If somebody in Norway made an observation, and I am trying to reproduce it here, how do I know that I'm observing the flies in the same way?"

"I had two biologists labeling my frames," remarks Perona postdoc Piotr Dollar, "and they only agreed on about 72 percent of them. And when I had the same person relabel some of them two months later, there was only about 80 percent agreement" with that person's previous set of observations. It wasn't merely a matter of changing one's mind on where an action started or stopped, either—"entire behaviors would be missing."

Since computational techniques for identifying dark objects on a bright background (or vice versa) are well established, as are methods for following those objects from one frame of a video to the next, Perona figured it might take about three weeks to write Anderson's and Dickinson's software. In fact, the first draft of each package took more than a year. Finding the flies and keeping track of them was the easy part; "the big difficulty was keeping the flies distinct when they were bumping into each other and overlapping," says Perona. It also proved harder than expected to make systems that would run reliably in the hands of people who were not computer scientists.

The first step was bringing the video-labeling process into the electronic age, with a captioning package not unlike the video-editing software you might have on your PC. "There was a commercial system available, but it wasn't terribly good," says Perona. "So we had to write one from scratch." As a fly movie plays in one window, another window displays a separate time track for each insect. Clicking on the time track allows you to mark the frame where, for example, the fly begins its lunge. "We would play the video at one-tenth speed, freeze it, click on a fly, and then choose a behavior from a dropdown menu," says Liming Wang, a grad student of Anderson's.

Hundreds of hours of meticulously annotated video would be fed to the computer, which could then teach itself what a lunge looked like by scanning the database, watching all the segments marked "lunge," and extracting some set of parameters common to all the examples. Once the computer had processed this "training set," the researchers would give it a fresh set of videos, the test set, for it to label on its own. Of course, the humans had to check its work, which meant that all of the test set's videos had to be labeled by hand as well. And then the process would repeat. Endlessly.

It took about three years of machine learning to get ready for prime time, but the finished products debuted last spring. Both software packages can be downloaded for free, and they are now in widespread use. The one for watching pairs of flies—the cage matches, if you will—is called CADABRA, for Caltech Automated *Drosophila* Aggression-Courtship Behavioral Recognition Algorithm. It was created by Heiko Dankert, a postdoc working with Perona and Anderson; Wang; and Anderson postdoc Eric Hoopfer. The other, which records and displays the meandering paths of large groups of flies in an open field, is named Ctrax—pronounced "See-tracks," get it? It was developed independently by postdoc Kristin Branson, working with Perona and Dickinson, and Dickinson's grad students Alice Robie (PhD '10) and John Bender (PhD '07). CADABRA and Ctrax were published in the April and June 2009 issues of *Nature Methods*, respectively. (Dickinson has since joined the faculty of the University of Washington in Seattle.)

CADABRA: MAKE LOVE, OR WAR

The behaviors that CADABRA automatically recognizes include three that are explicitly aggressive: the "tussle," in which flies sumo-wrestle by facing each other, gripping one another with their forelegs, and struggling to displace their opponent; the "lunge," in which one fly rises up on its hind legs and pounces on its opponent; and the "wing threat," in which both wings are extended perpendicular to the body and then tilted up about 45 degrees, presumably to make the fly appear bigger and more intimidating. Three others are courtship-related: "circling," in which the male fly walks sideways around the object of his desire, facing head-in; "wing extension," in which our suitor woos his intended by vibrating an outstretched wing in a courtship song; and "copulation," over which we shall draw a discreet curtain. The final one, "chasing," can lead to either sex or violence, depending on the circumstances.

CADABRA scrolls through the video frame by frame, locating and identifying each fly. (Females are bigger than males, making it easy to keep straight who's who; when male flies are paired up, one of them gets a

Figuring out which way the fly is facing provides critical information, as butting heads and bumping butts are likely to lead to very different outcomes.

drop of white paint on his back.) Then the software analyzes each fly's pose—creating a mathematical representation of its body language. The process begins by breaking each fly down into three ovals: one for the body and one for each wing. The wings are a different shade than the body, and thus easily distinguishable. Figuring out which way the fly is facing provides critical information, as butting heads and bumping butts are likely to lead to very different outcomes. Fortunately, fly heads are quite shiny due to their reflective, waxy cuticle, so the computer divides the body oval in two and labels the brighter half the head. Then the system calculates a set of 25 parameters, starting with the fly's position, velocity, and direction of travel and going into such details as the body's apparent length, its angle of orientation (which is not necessarily the same direction it is moving), and the angles of the wings to the body.

"Heiko, Pietro, and I watched a whole bunch of movies and made a big list of what parameters the system needed to look for," says Wang. "In a lunge, for example, when the fly rears up to strike, it gets smaller as seen from above. When it lunges, the head suddenly accelerates. And before it lunges, it often stops, to sort of gather itself. It follows its opponent, then pauses, then lunges." It usually took 10 to 15 of the 25 parameters to describe each behavior, Wang says. If the action was simple enough, the biologists could even define it explicitly—for example, in a wing extension, the wing had to be outstretched between 60 and 90 degrees from the body for at least one second. This precision was alluring, but the more complex activities eluded such easy encapsulation. All the biologists could do then was label the videos and leave the machine to figure things out for itself.

CADABRA's analytical power comes from its ability to compile videos into "stacks" that can be compared to other stacks. A stack could be a "heat map" that shows where something tended to happen; for example, fights were most likely to break out at a food source placed in the center of the arena, causing that area to glow red. Behavioral differences started to pop out immediately. Males from a strain of flies called Cha-Tra, whose brains have been "feminized" by messing with the genes that control masculine development, spent an awful lot of time chasing each other around the arena's periphery. Pairs of unmutated, or "wild-type," males were more venturesome, and were as likely to chase each other out in the middle of the arena as they were to hug the walls. Timelines can be stacked just as easily. "Flies fight a lot in the first 10 minutes," Wang says. "After that, they seem to get to know each other. Or they get bored. We don't know."

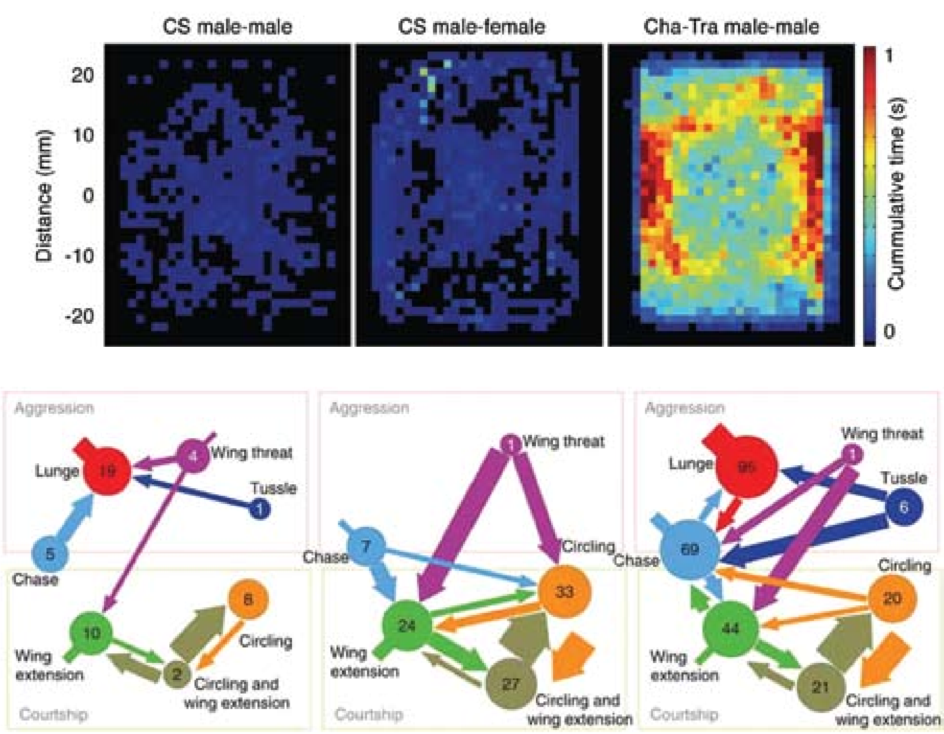

Top: A trio of CADABRA heat maps. The hotter the color, the more time pairs of flies in that location spent chasing each other. CS (for Canton-S) flies are a standard strain of unmutated flies. They went their own way most of the time, but their few chases were fairly evenly distributed throughout the arena. Pairs of Cha-Tra males were much more likely to chase each other along an arena wall. Bottom: CADABRA ethograms for the same pairs of flies. The size of each circle, and the number in it, reflects how often each action occurred. The relative widths of the connecting arrows show what was most likely to happen next. (In order to qualify as related, the two actions had to be separated by 10 seconds or less.) The stubs show the probability of the same action being repeated: thus, if a Cha-Tra fly lunged, it was likely to lunge again. Twenty pairs of CS males, 24 male-female pairs, and 10 pairs of Cha-Tra males were recorded in this set of videos.

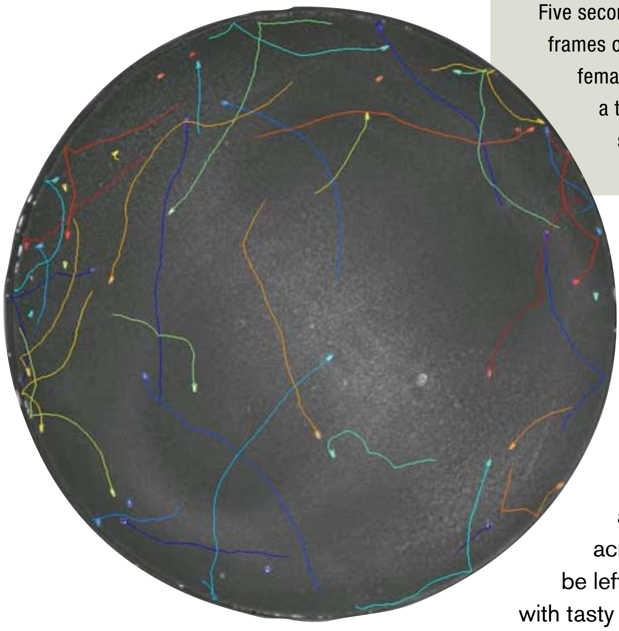

Five seconds' worth of Ctrax data—100 frames of video—for 50 wild-type female flies. Each fly is marked by a tiny triangle that draws a line showing where it's been.

Most importantly, CADABRA can create the behavioral equivalent of a traffic-pattern diagram. Called an ethogram, it maps how actions flow into one another. For example, when wild-type males were paired up, a chase was usually followed by a hostile lunge. But put two Cha-Tra males together, and they were as likely to become lovers as fighters, with a chase leading to an amorous advance in the form of a wing extension as often as it did to a lunge. Such analyses might reveal whether aggression and courtship are at opposite ends of one continuum, controlled by a dimmer switch—the level of one or two critical proteins, perhaps—or whether they're really two different states of mind, if flies can be said to have minds, controlled by two independent neural circuits.

With CADABRA up and running, the Anderson lab has shifted into high gear, screening 100 different strains of flies per week, and recording a dozen pairs of flies per strain. Fully annotating seven actions per video would have soaked up more than 80 person-hours per strain—two entire work weeks for some poor grad student. CADABRA does the entire analysis in a few minutes.

CTRAX: FACES IN THE CROWD

Ctrax, created for the Dickinson lab, monitors groups of flies wandering as the whim takes them in an arena 10 inches in diameter—scaled to the length of a fly's body, that's about a mile across. This vast landscape can be left empty, or it can be strewn with tasty tidbits at which to gather (or fight), obstacles to explore or navigate around, pinnacles to climb, or anything else the researchers can come up with. "It is a bit of a fishing expedition," Perona says. "If the machine can look at kilometers of footage and detect some regularities, it may be able to formulate a hypothesis that a biologist hadn't yet made." And there may be behaviors that are so rare humans might not see them, or that take so long to play out we would not realize that they are happening.

Postdoc Kristin Branson faced an enormous challenge when writing Ctrax. How could she keep each fly's identity straight without having to attach a physical marker to the insect, as the best commercially available package did? Other systems that didn't use markers were prone to confusion when the flies got close together, and would often misassign their IDs when they parted ways again. Ctrax minimizes identity theft through frame-by-frame comparisons. It first looks at each frame in isolation and tries to locate all the flies it can by using what Perona postdoc Michael Maire calls a "blob detector." Then, working on the assumption that each fly won't have moved much, it compares consecutive frames to see whether a blob appears in roughly the same spot and is moving in a smooth path. If the blob is too big, the computer subdivides it by fitting fly-sized ellipses to contiguous groups of pixels. (Even if the flies are climbing over one another, there's usually a gap of a few pixels' width between some portion of their overlapping silhouettes.)

Backtracking along each ellipse's path tells Ctrax which end of each blob to call the head—flies are more likely to be walking forward than backward. Flies don't usually enter or leave the arena, so if a track suddenly vanishes during the rewind, the computer looks around for other tracks. If a nearby track dead-ends, for example, perhaps it and the vanishing track should be spliced together. Or perhaps a track forks, and the vanished track can supply the missing leg. The system can get stumped if, for example, a fly rears up and appears foreshortened, or makes an abrupt move in a radically different direction. Ctrax then calls for help, asking a human to look at the video and fix things by hand. "This happens once every fly-hour or so," says Perona.

Once the processing is complete, Ctrax displays each fly as a thin triangle—the pointy end being the head—trailing a colored line that traces its path. The line usually just shows the last little bit of the fly's history; otherwise, the arena quickly fills with a rainbow of scribbles resembling the work of a bored five-year-old with a fresh box of Crayolas.

The arena quickly fills with a rainbow of scribbles resembling the work of a bored five-year-old with a fresh box of Crayolas.

Dickinson's lab trained Ctrax to recognize walking, stopping, turning sharply, backing up, jumping, chasing, touching, and crab-walking, in which the flies move sideways. Then they gave the system a workout, using 17 groups of 20 flies each: all-male contingents; all-female ones; 50-50 mixes; and batches of male flies with the *fru* mutation, which controls male courtship and mating. By creating an ethogram for each individual fly, Ctrax could automatically determine its gender (and, in the case of the *fru* flies, its phenotype) with better than 95 percent accuracy. Male flies proved to have little sense of personal space, routinely approaching other flies of either sex until little more than a body length separated them. Females were more retiring, preferring to keep at least two body lengths between themselves and others. And *fru* flies were far more likely to back away from an encounter than were wild-type males.

Each wild-type fly displayed its own self-consistent neural programming. Some flies kept close to the walls, while others ventured out into the center more often. Each fly had a preferred walking speed—the zippiest striding twice as fast as the pokiest—and a favorite duration for its strolls. Even the percentage of time a fly spent moving around versus standing still was an individual trait.

Ctrax's ultimate validation came when it flagged large numbers of track segments that defied classification. When the humans reviewed them, several new behavior types were identified: T-stops, X-stops, jousts, drag races, and even games of chicken.

FROM SLEDGEHAMMERS TO THERMOSTATS

Seymour Benzer's groundbreaking work back in the 1960s had used a sledgehammer approach, zapping the flies with X rays or dosing them with chemicals to induce wholesale mutations. The screening process was quick, thanks to ingenious methods his lab developed to collect the flies displaying some desired behavior. The time-consuming part of the job came afterward: inventorying the mutations each fly carried and then trying to figure out which one actually made the difference.

Nowadays, Anderson turns neurons on or off as easily as flicking a light switch—or, rather, adjusting a thermostat, thanks to a nifty piece of molecular biology developed at Brandeis University by Paul Garrity (PhD '93). The procedure exploits a temperature-sensitive ion-channel protein called TrpA1, for Transient Receptor Protein A1, which is normally found in heat-sensing neurons that help the fly stay in its comfort zone. Flies like it a bit on the chilly side, so TrpA1's ion channel stays squeezed tightly shut at 22°C, but unclenches at 27°C—the equivalent of taking the flies' cage from an air-conditioned lab out into a nice summer afternoon. When the channel is closed, the neuron can't fire. When the channel opens, the neuron goes off, telling the fly to start looking for a cooler place to hang out.

An extra copy of the *TrpA1* gene can be inserted into a fly's DNA globally and selectively activated in some specific set of neurons that are not normally heat-sensitive. The details are complicated, but depending on the type of neuron you choose, you can tweak a set containing only a few tens of cells. Other sets might have a few hundred, or perhaps a thousand neurons—still just a few percent of the fly's nonvisual brain cells.

The activated neurons proceed to sprout lots and lots of extra copies of the ion channel on their surfaces. At 22°C, the modified neurons go about their business as usual, firing at their normal rate whenever they're supposed to fire. When the temperature rises to 27°C, they still fire whenever they're supposed to fire. But because of all the extra ion channels, when they *do* fire, they just go nuts.

Traditional aggression screens have focused on genes, not neurons. These studies "knock out," or inactivate, a gene to see whether the resulting flies are less (or more) quarrelsome. But *over*activating a neuron and looking for flies with hair-trigger tempers has its advantages. "There are a lot of uninteresting, low-level ways to break a complicated behavior," says Anderson. "If the fly's legs don't work, they're not going be aggressive. But if you can specifically *enhance* aggressive behavior, there's much less likely to be a trivial or uninteresting explanation for it." And if the flies mellow out when retested at room temperature, "it tells us the increased aggression really is due to activating this particular subset of neurons, and not because those flies just happened to have had a bad day and were in a bad mood and fought a lot."

Studying neurons instead of genes has another benefit. If many genes are involved, which is likely, each one might make a small, subtle contribution. "But a neuron reflects the combined activities of whatever 12,000 to 15,000 genes are turned on in that cell," says Anderson. A pilot study of a couple

The colored boxes in this two-minute trace (a) of a male fly's movements through a coed crowd show seven behaviors that Ctrax labeled. In a close-up look at those boxed portions of the track (b), the triangles mark the fly's position in each frame. The blue and red triangles are the beginning and end points Ctrax assigned to the action; in the "walk" example, only the beginning point is shown. Gray triangles indicate the presence of a second fly. Plotting these actions against time (c) creates a visual summary of the fly's activities.

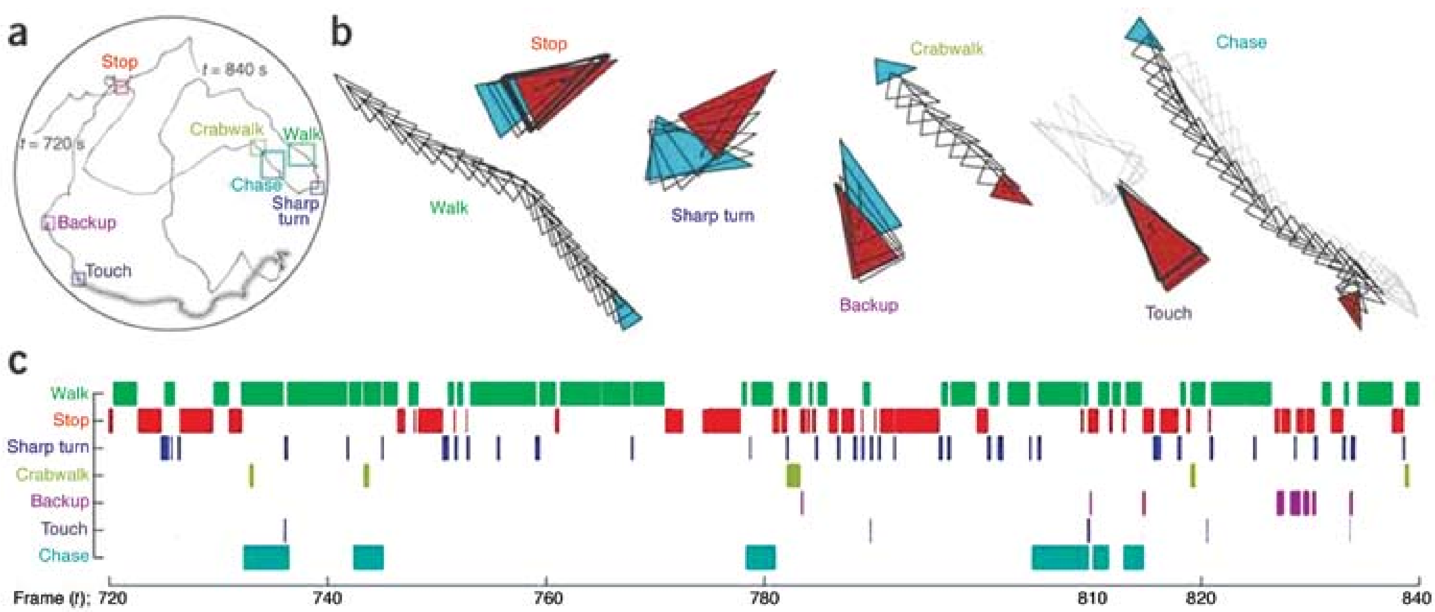

Adapted from Branson, et al., *Nature Methods* 6: 6 (2009) © Macmillan Publishers Ltd.

*Over*activating a neuron and looking for flies with hair-trigger tempers has its advantages. "There are lots of uninteresting, low-level ways to break a complicated behavior," says Anderson. "But if you can specifically *enhance* aggressive behavior, there's much less likely to be a trivial or uninteresting explanation for it."

hundred strains of flies showed that in the few that were extra feisty, "the increase is *very* big and *very* dramatic—by an order of magnitude. You don't need fancy statistics to see it."

Bigger studies are on the horizon. Anderson's postdoc Eric Hoopfer is now at the Howard Hughes Medical Institute's Janelia Farm Research Campus in Ashburn, Virginia, home to a collection of flies with overabundant TrpA1 channels in about 4,000 different sets of neurons. Hoopfer has adapted CADABRA to run on Janelia's supercomputer, and he plans to examine videos of every last one of those strains—the most comprehensive aggression screen ever attempted. "We're not making any assumptions about what *part* of the brain is involved in aggression," says Anderson, "or what *kind* of neurons are involved in aggression. We just test as many lines as possible. If some neurons keep showing up over and over again, or maybe look like they might be connected to each other, we can use this information to try and piece the circuit together. This software has really enabled an approach that could not previously have been undertaken."

OF MICE AND MEN

All through this past five years of fly work, Anderson never gave up on mice. "I like the idea of studying the same behavior in two evolutionarily very different species," he says. "Despite their obvious differences in brain structure, are there are some general principles that underlie the organization of aggression circuitry?"

We may be on the road to finding out. Perona postdoc Piotr Dollar is working on a generalized version of CADABRA that, alas, does not yet have a catchy acronym. Mice, being fluffy and flexible, are much harder for a computer to discern. They can curl up in a ball, or stretch out, or hunch over and scratch behind an ear. To make matters worse, this nearly infinite variety of "looks" has to be extracted from a textured background of wood shavings that registers in the same shades of gray.

Dollar has tackled such problems before. He used to work on the Caltech Pedestrian Detector, a program intended for cameras that could, for example, be mounted on airport shuttle buses. Such a system would be intended to alert the driver that someone is about to step out from behind a parked car and into the crosswalk ahead, and so it would have to deal with partially obscured bodies and blotty backgrounds all the time.

Says Perona, "Piotr figured out a very clever way, which is true progress in machine vision, for detecting the mouse." Dollar's system uses a collection of "weak" feature detectors, whose outputs are collected to render a verdict, Perona explains, and "his detectors are designed around sophisticated visual measurements that nobody had managed to use in practice before. It turns out that his method extends to almost any animal, and you can train it very easily." Says Dollar, "You just draw a little circle around the animal of interest in a few hundred frames, and then you also give the system a bunch of negative examples, which could be pictures of anything—you don't care."

Dollar is collaborating with Andrew Steele, a Broad Senior Research Fellow in Brain Circuitry, to test the system. "Getting something to work robustly in different lab environments is what makes the thing click," says Steele. "What separates a really great computer-vision person from someone who's merely good is that they give you something you can actually *use*. Because a lot of these papers that people publish only work with one database. They tune their algorithm to work really well on one type of video, in one lighting situation. And then it's not very useful for the end user like me."

Figuring out how to tell a mouse's head from its tail is going to take a while, however, and training the system for behavior recognition is still very much a work in progress. Anderson estimates that his lab—including a bunch of undergraduates on work-study—has invested 1,500 person-hours in annotating mouse videos.

But the payoff will be enormous. Says Anderson, "There are projects going on around the world—at the Sanger Institute in England, and at other sites—to generate a complete library of mutant mice in which each mouse has one, and only one, of its 20,000 genes inactivated. With a computer program analogous to CADABRA for mice, it would then be possible, in theory, to screen through all 20,000 mutants to see which ones have the biggest influence on aggressive behavior. And you would know in advance which gene was knocked out, which would be a *huge* advantage."

Brain function is as much about chemistry as circuitry, so Anderson's lab is also exploring the effects of pheromones and other chemical messengers such as dopamine. He's been doing this all along with the flies, but the mouse work could be adapted for what the biomed biz calls "translational science"—screening drugs to treat impulsive violence, for example. "We're not equipped to do it," Anderson says, "but it's something that the pharmaceutical industry might be very interested in doing." In the longer term, he adds, "current treatments for psychiatric disorders are very suboptimal. We have little understanding of what goes on in, for example, the brain of a depressed person, or how depression alters brain function. We need to understand the construction and function of the normal circuits that process emotional behaviors in order to understand how that function can become abnormal." E&S

Map by Veronica Olazabal for *Who's Watching? Video Camera Surveillance in New York City and the Need for Public Oversight* by the New York Civil Liberties Union, Fall 2006.

BIG BROTHER (AND EVERYONE ELSE) IS WATCHING YOU

A block-by-block survey of lower Manhattan conducted by the New York Civil Liberties Union in 2006 counted 4,176 security cameras between Battery Park and Fourteenth Street—and those were just the ones visible from the sidewalk. The city of London is said to have half a million of them, in public and private networks. "Right now, the cameras are mostly just videotaping, and storing the video," says Perona. "But in principle, you could network these cameras and automatically analyze what they are seeing."

Such automatic surveillance could discern forms of suspicious behavior too subtle for a human watchman to pick up at a distance. "Suppose you have a terrorist wearing a 30-kilogram explosive vest," says Perona. "He might walk in a slightly different way, because of the extra weight on the upper body." Other warning signs might slowly emerge over days or even weeks, eluding all but the most acute observers. "Say there's somebody sitting on a bench, pretending to be reading a newspaper, but taking mental notes for a future attack. If you observe the scene, there is nothing wrong with it. But if you see the same person the next day hanging around a phone booth, making a three-hour phone call, and on the third day, there they are again, sitting somewhere else, you will want to ask them what they are up to."

"These systems are coming," says Perona, "and they will be used. And the public needs to help regulate them in the proper way." —*DS* E&S

***David Anderson* is the Benzer Professor of Biology and an Investigator at the Howard Hughes Medical Institute. He earned his PhD at Rockefeller University in 1983 and came to Caltech as an assistant professor in 1986.**

***Pietro Perona* is the Puckett Professor of Electrical Engineering. He earned his PhD at Berkeley in 1990 and moved south to Caltech the following year.**

***Michael Dickinson* earned his PhD at the University of Washington in 1989. He arrived at Caltech as a visiting associate in 2001 and became a full professor here the following year. He was the Zarem Professor of Bioengineering at Caltech from 2003 to 2010.**

On November 18, Anderson, Dickinson, and Caltech neurobiologist Christof Koch were named to the inaugural group of Allen Distinguished Investigators by the Paul G. Allen Family Foundation—seven scientists "working on some of the most exciting research in biology and neurology," according to Microsoft cofounder Paul Allen. (For more on Koch's work, see page 14.)

The CADABRA project was funded by the National Science Foundation, the National Institutes of Health, and an Alexander von Humboldt Foundation grant to Heiko Dankert. The Ctrax project was supported by the National Institutes of Health.