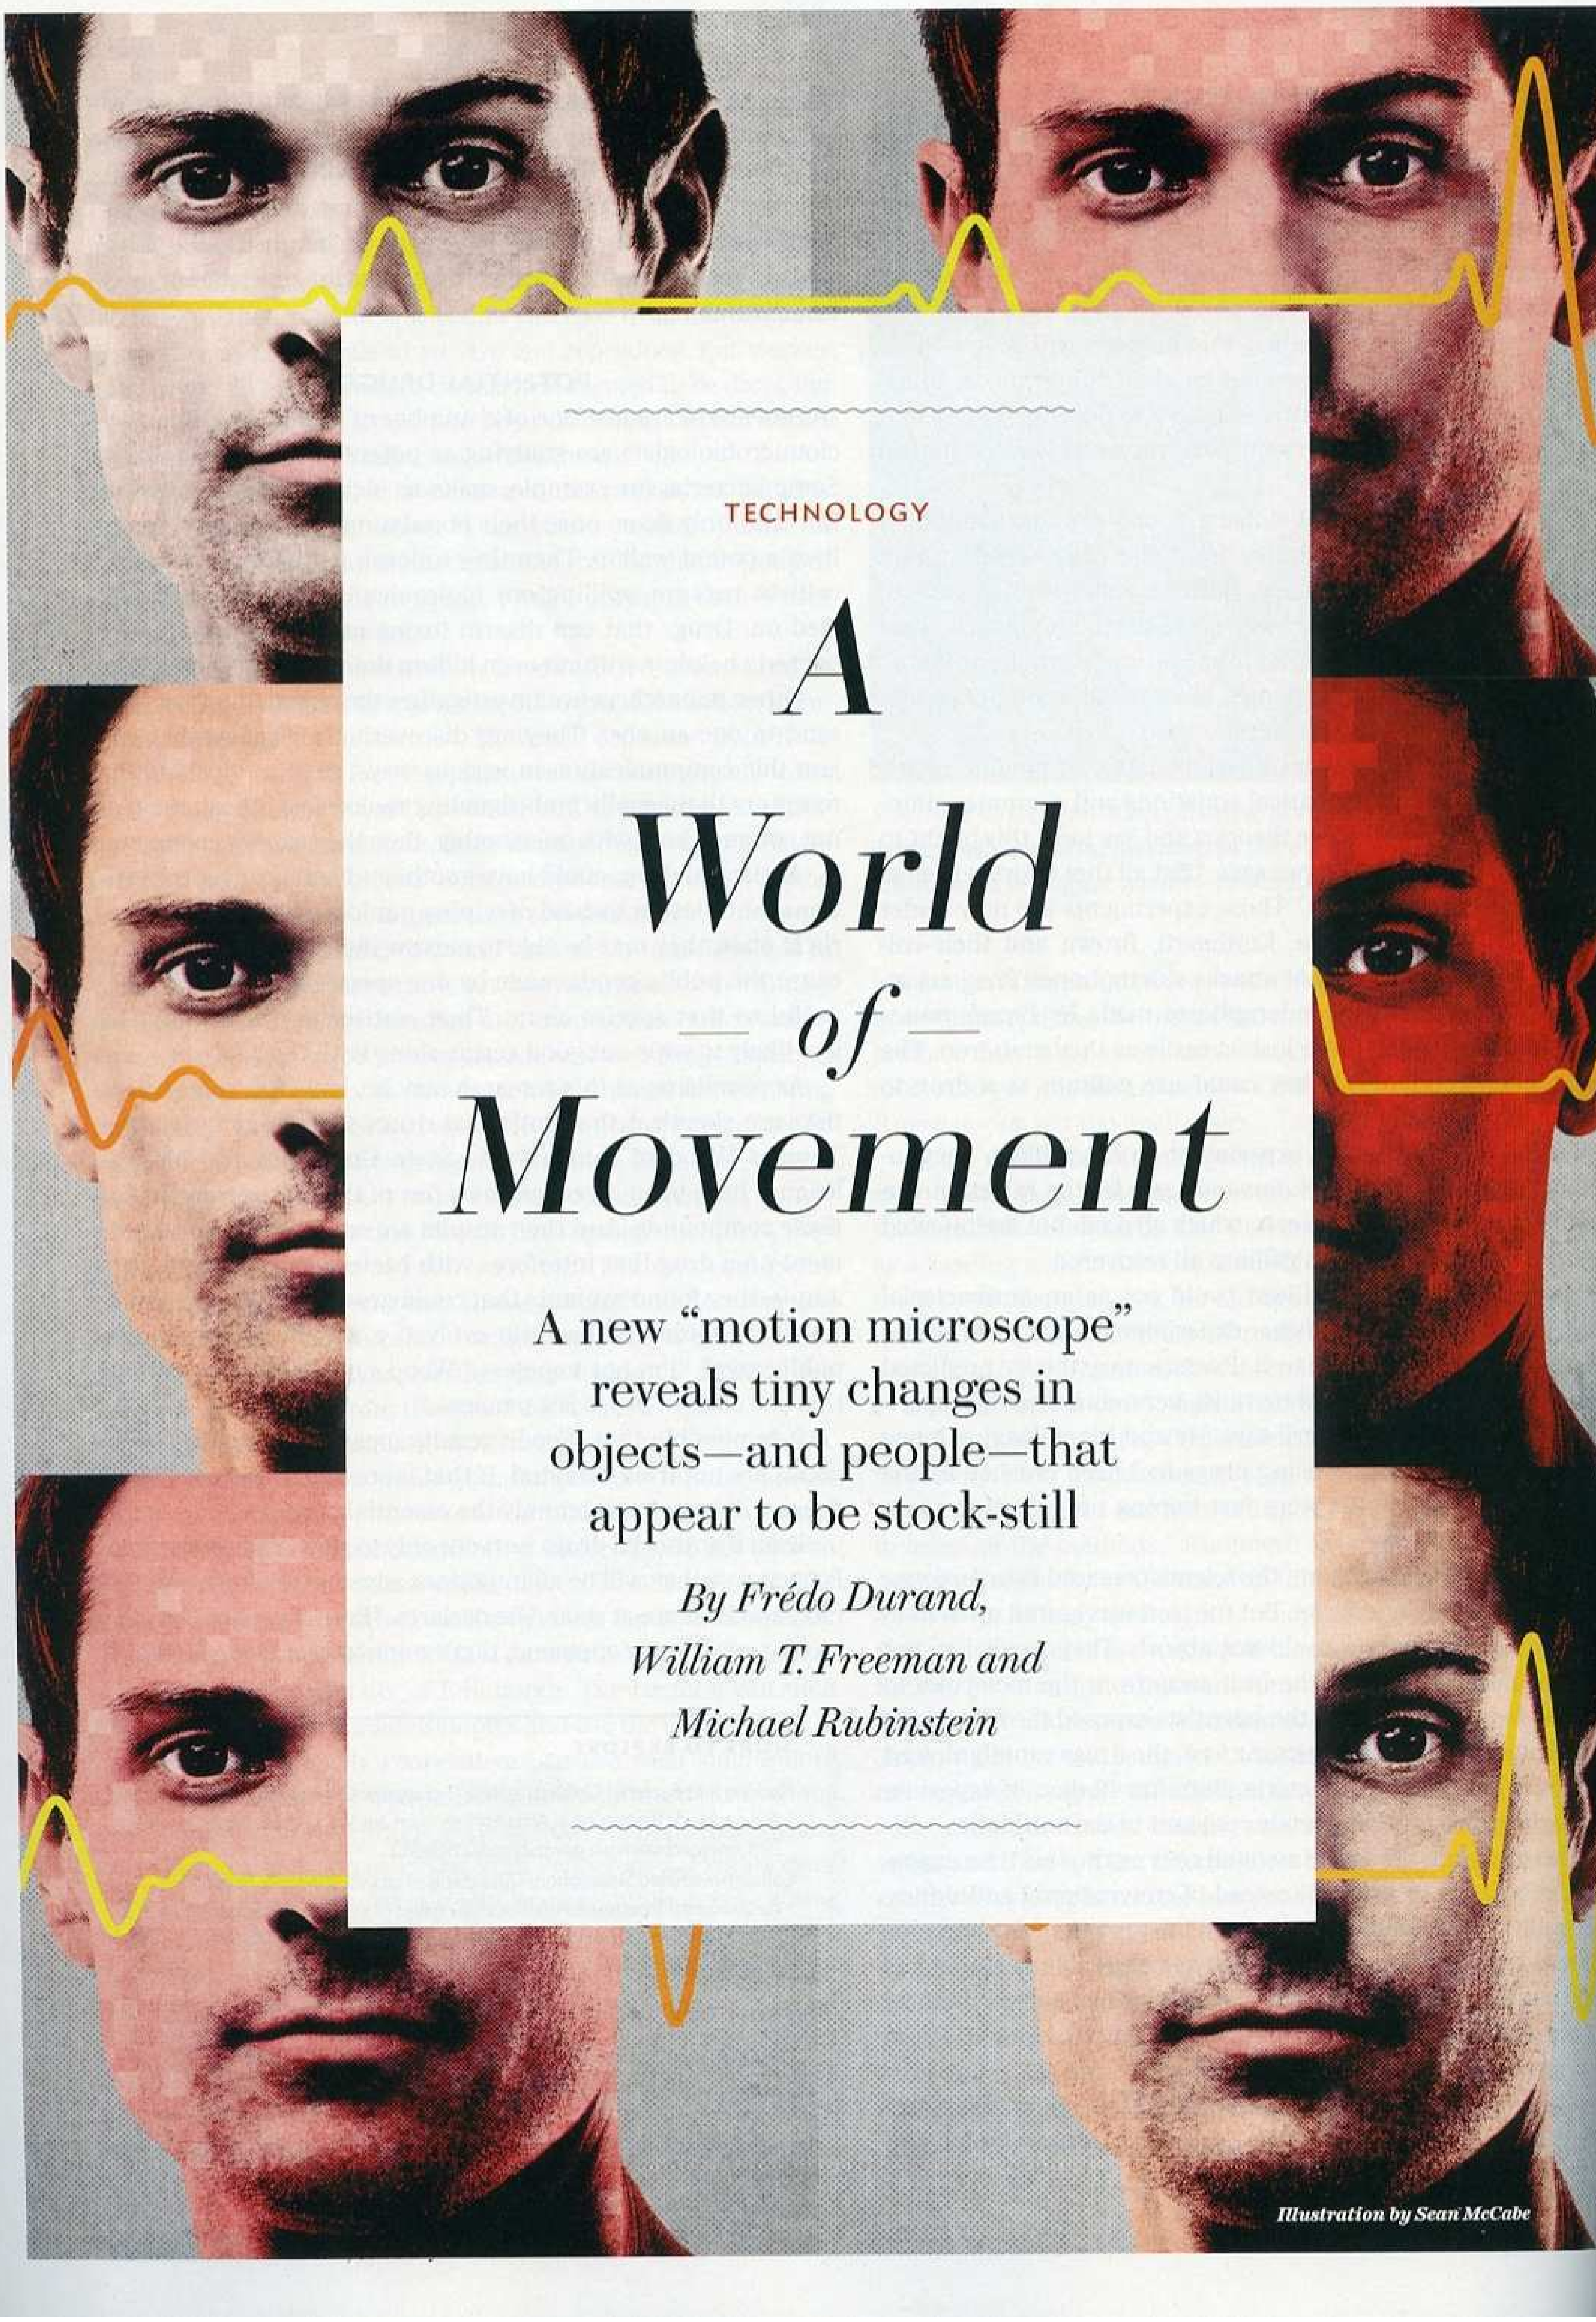

TECHNOLOGY

A World of Movement

A new "motion microscope" reveals tiny changes in objects—and people—that appear to be stock-still

By Frédo Durand, William T. Freeman and Michael Rubinstein

Illustration by Sean McCabe

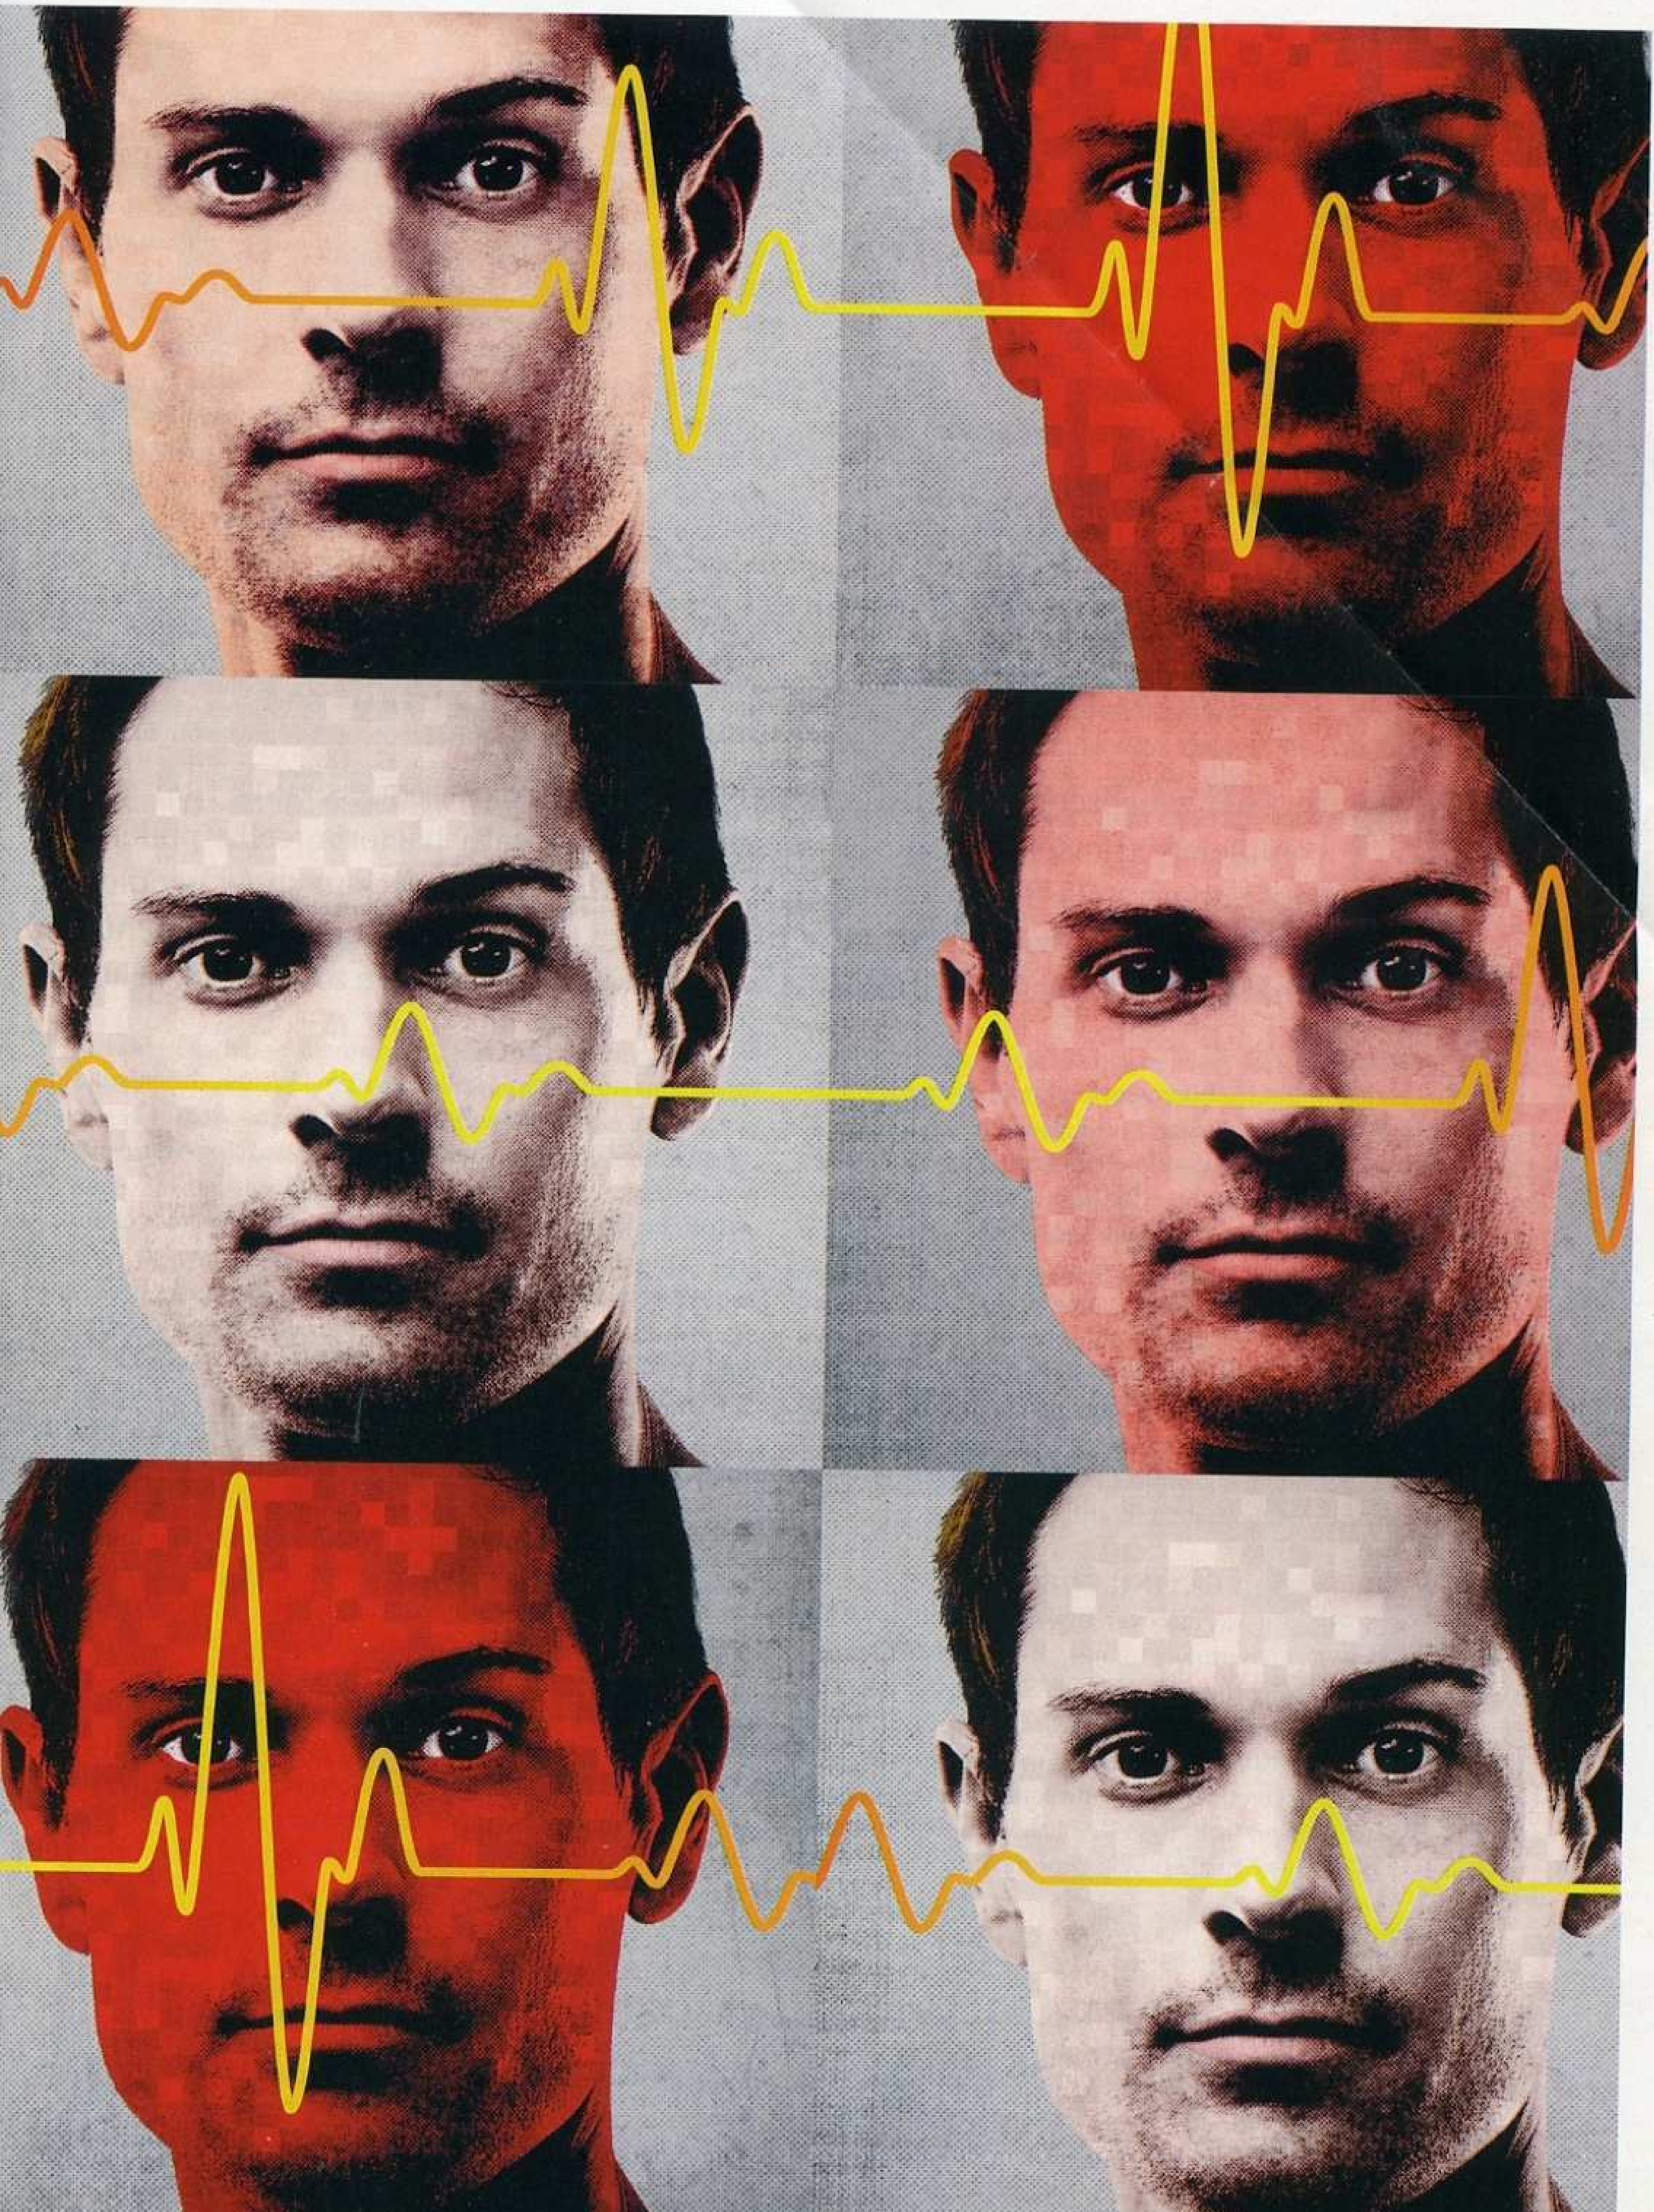

Frédo Durand works on computational photography as a professor of electrical engineering and computer science at the Massachusetts Institute of Technology.

William T. Freeman, a professor of electrical engineering and computer science at M.I.T., studies how machine learning can be applied to computer vision.

Michael Rubinstein is a research scientist at Google, working on computer vision. His motion microscope work was done while he was at Microsoft Research and M.I.T.

THE FIRST MICROSCOPES, IN THE 1500S AND 1600S, TRANSFORMED GLASS PANES THAT LOOKED completely transparent into a universe teeming with bacteria, cells, pollen and intricate crystals. These visionary aids were the first devices to show people that there were cells within a drop of blood. Since then, microscopes have opened up other invisible worlds for scientists, going within cells or down to the scale of atoms.

We believe a new kind of microscope is about to unveil another fascinating new world: a world of motion and color change too minute for the eye to catch. Blood pulsing through one's face makes it redder and then lighter, the wind can cause construction cranes to sway by a tiny amount, and a baby's breathing is often too subtle too be seen. These movements are almost unimaginably small, yet their importance looms large. They can reveal the state of our health or the vibrations of a crucial machine about to fail. With our students and collaborators, we have developed what we call a motion microscope, a tool that couples a video camera with specialized computation. Together they amplify movements in people and objects that seem, to the naked eye, to be standing absolutely still.

CALCULATING COLOR

OUR MOTION MICROSCOPE was discovered serendipitously. We had been working on a video project to measure tiny color changes, too small to be seen by the unaided eye. Scientists Ming-Zher Poh, Daniel McDuff and Rosalind W. Picard of the M.I.T. Media Lab had shown, in 2010, that they could use a video camera to measure a pulse by detecting minuscule color variations caused by blood flowing to and from the face in rhythm with the beats of the heart. (They have turned the technique into a pulse-measuring smartphone app called Cardiio.) We felt the calculations were tricky and more complex than they needed to be, involving advanced linear algebra. We began searching for a simpler way to carry out the process.

The main challenge is the low degree of the color change in any individual video pixel caused by blood flow—it varies by just 0.2 percent over the course of a pulse. Unfortunately, camera sensors do not record exact values and always contain random noise, typically higher than 0.2 percent. This noise vastly overshadows the variation in redness.

In our search for a simpler route, we, along with our then student Hao-Yu Wu, researcher John Guttag of the Massachusetts Institute of Technology and Eugene Shih, then at Quanta Research Cambridge, decided to replace the number representing the color of each pixel with an average of all nearby pixels. This method dramatically reduced the noise because these random fluctuations tend to cancel one another out within a large enough pixel group. We also filtered out color changes that occurred over a longer or shorter period than the range typical of the resting pulse for adults.

Our simple approach proved successful at converting the pixel changes into the number of beats per minute. But these color changes were invisible to us, and we wanted to see what they

IN BRIEF

Although they appear to be absolutely still, objects and people move in ways imperceptible to normal vision. These movements can be as small, and as important, as a baby's breaths.

By amplifying color changes in video pixels as they change moment by moment, researchers have created a "motion microscope" that makes these small motions very visible.

These magnified movements can show crucial health indicators such as changes in pulse or blood flow or looming safety problems such as abnormal vibrations in heavy machinery.

Frédo Durand works on computational photography as a professor of electrical engineering and computer science at the Massachusetts Institute of Technology.

William T. Freeman, a professor of electrical engineering and computer science at M.I.T., studies how machine learning can be applied to computer vision.

Michael Rubinstein is a research scientist at Google, working on computer vision. His motion microscope work was done while he was at Microsoft Research and M.I.T.

looked like. By using these calculations to compute changes in redness at each pixel in a video over time and then amplifying them by a factor of 100, we were able to clearly see the face of an adult man getting redder every time his heart beat.

This technique also works for babies. In a test on newborns, performed with physicians Donna Brezinski and Karen McAlmon, both then at Winchester Hospital in Massachusetts, we shot a video with a regular digital camera. After amplification, we found the pulse shown by the video and a traditional pulse meter attached to a tiny finger matched closely. This observation raises the possibility of measuring a pulse without contact, which is important for fragile premature neonates because touching such infants can be harmful. For adults, in the future, these visualizations may help reveal abnormalities in blood flow that could have health implications, such as asymmetries in circulation between the left and right sides of the body.

NOT SO STILL LIFE

OUR VIDEOS, however, presented us with a puzzle. To simplify the color processing, we had asked the adults in front of our cameras to hold very still, and their heads really looked motionless in the original videos. But as we watched the color-amplified results, we noticed that their heads were moving. Our technique seemed to enhance not only color changes but also tiny motions.

In earlier work with other colleagues, we had created videos that amplified small movements. But that involved specialized software that computed motion directions—vectors—for each pixel at each point and moved them around to new locations. It turned out to be complicated and prone to errors. We were astounded that our new approach could achieve a similar effect with a simple computation and without calculating any tricky motion vectors.

Why did bigger color changes also magnify these small motions? To find out, we had to review how movement in a video results in local color changes. Imagine an object like a ball that is lit from the right, making the right side of the ball bright and the left side dark. If the ball is flying from left to right across a video screen, a pixel at one fixed location on the screen will get darker and darker as time ticks on because it depicts points farther and farther to the left on the ball. The variation depends on how fast the ball is moving and how sharp the brightness transition is between the left and right sides of the ball, the so-

BASICS

Turning Color Changes into Motion

In a video, each pixel represents a point on an object, such as a leaf or branch on a tree. As time ticks ahead, the color in that pixel changes as the leaf moves, even a tiny bit, because the leaf moves in relation to the light hitting it. A computer program that amplifies the color variation from one frame to the next also exaggerates the tiny motions so they become obvious to the naked eye. Researchers can isolate and amplify one particular time frequency—say, the speed of leaves shaking—and the result will make leaf movement stand out against the rest of the tree.

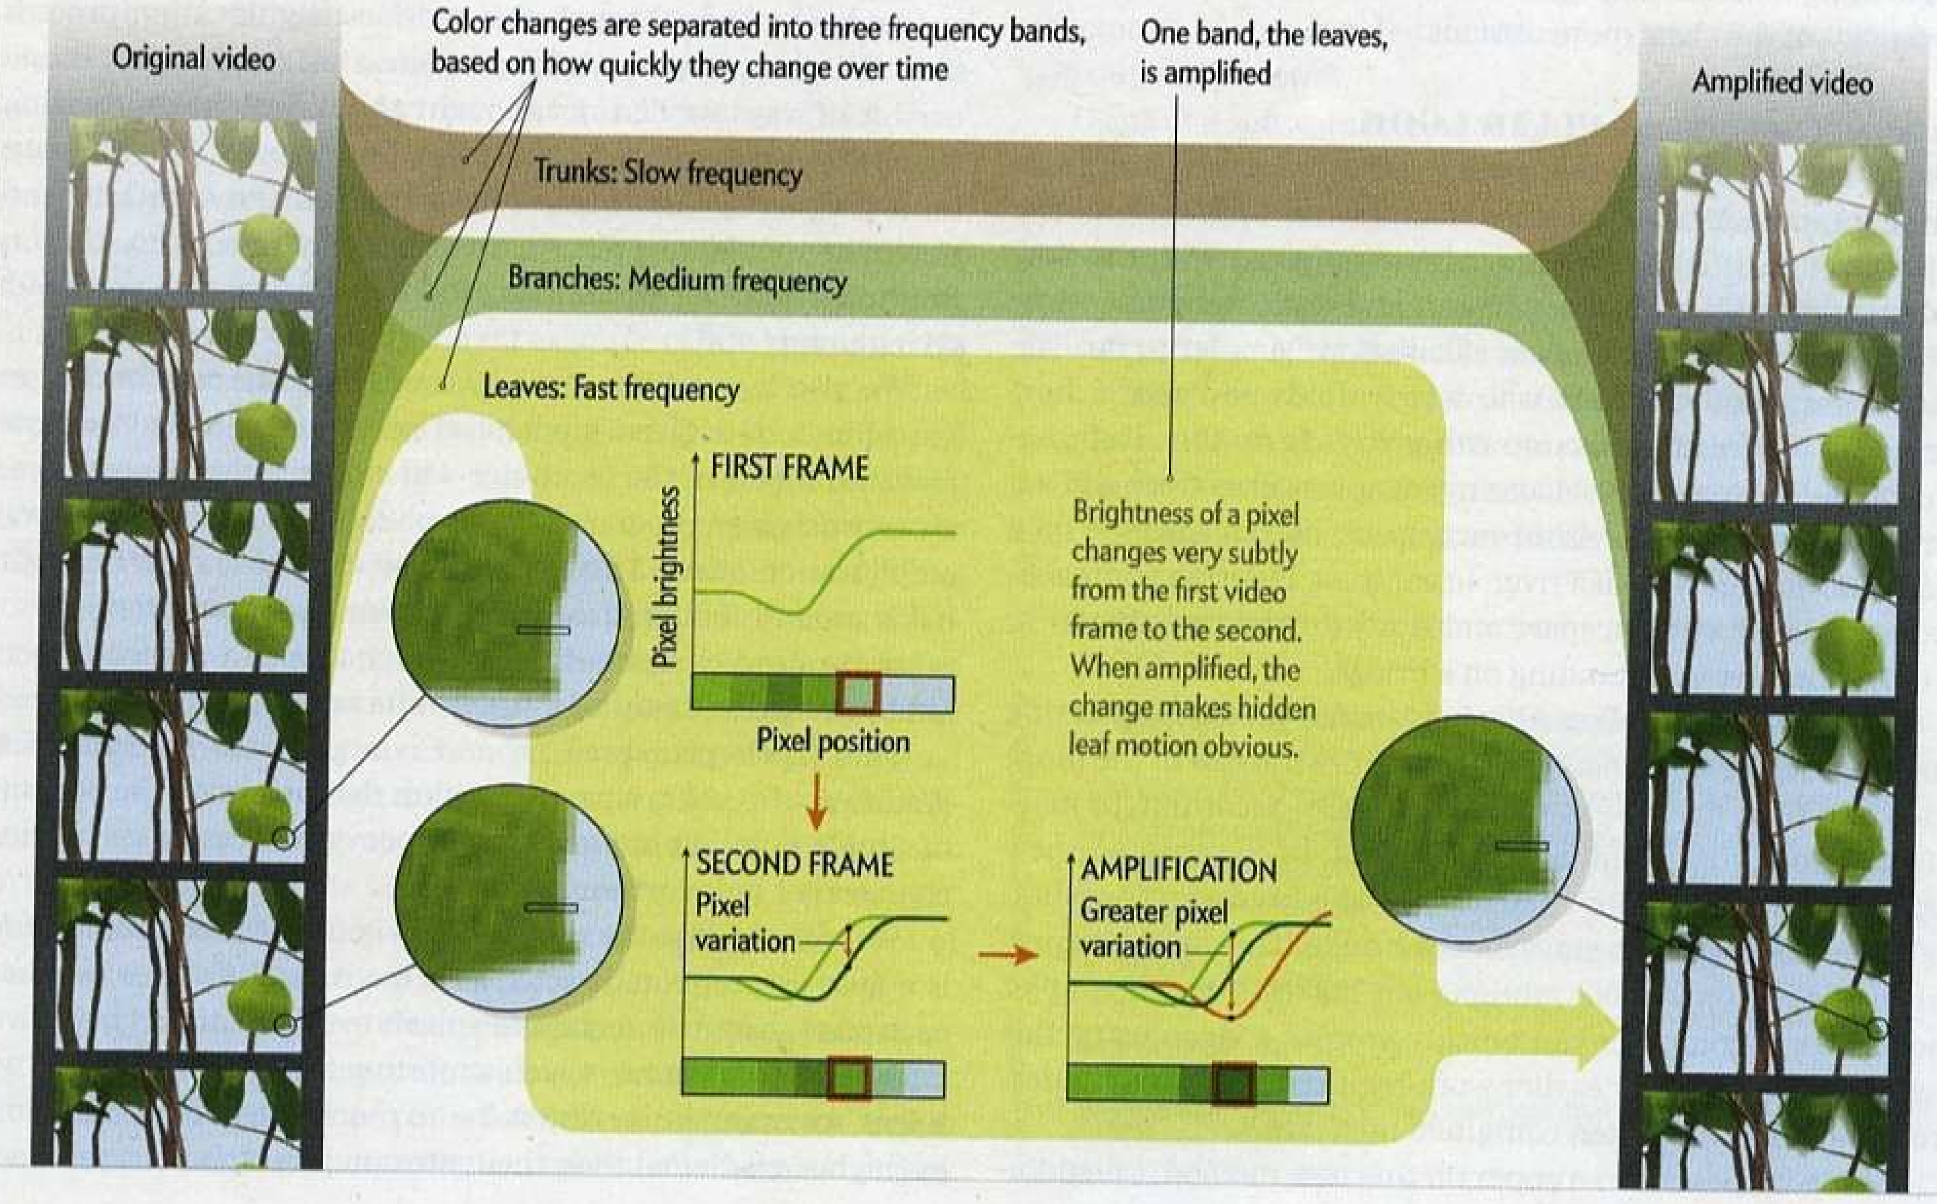

Brightness of a pixel changes very subtly from the first video frame to the second. When amplified, the change makes hidden leaf motion obvious.

Illustration by Jen Christiansen

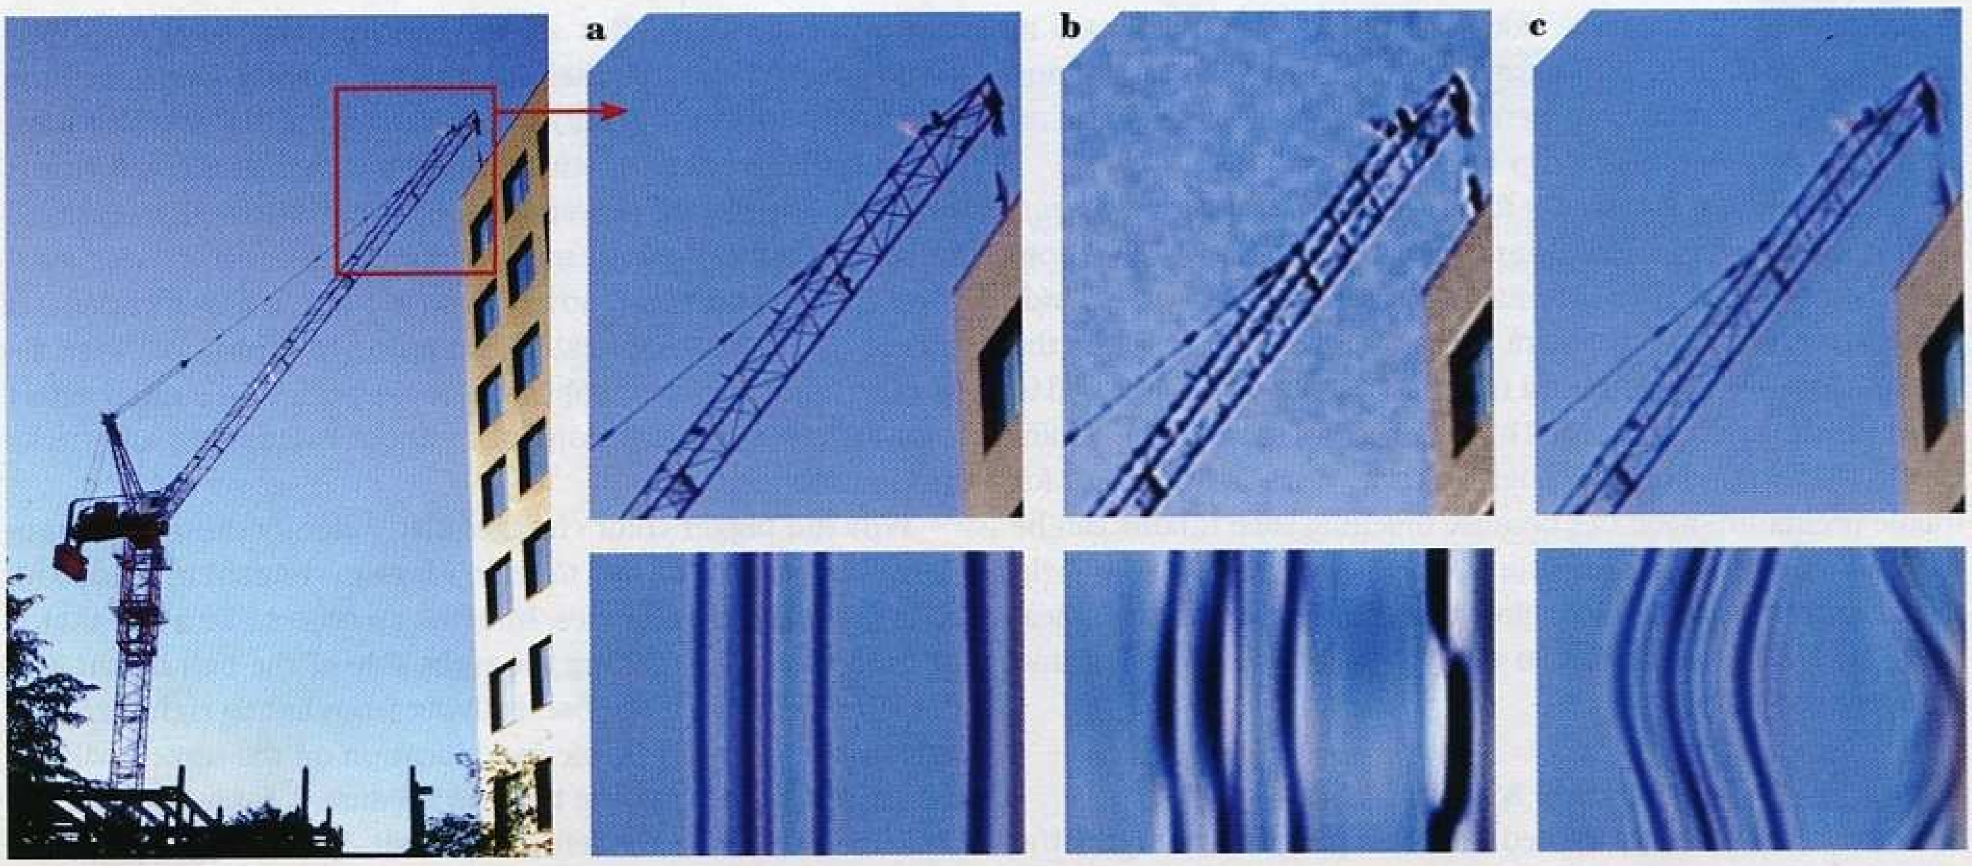

MOVING TARGET: This construction crane seems motionless (*a*). Amplifying video color changes reveals swaying, but pixels look jagged (*b*). A computer smooths pixel transitions, showing motion (*c*). The bottom images show one crane feature moving over time.

COURTESY OF NEAL WADHWA Massachusetts Institute of Technology

called color gradient. Mathematically, we can say that the change of a pixel's color over time is the product of the speed of the object multiplied by this color gradient.

Our algorithm, of course, does not know about speed or color gradients. Nevertheless, because it amplifies the color change at any particular point as the ball moves a fraction of an inch to the right, it also amplifies that fractional motion of the ball for your eyes to see. In a similar way, the colors of pixels representing specific points on a baby's chest will change as the baby breathes, and making the color change more dramatic also makes the tiny movement of the chest more obvious.

A FLUID LOOK

THE DIFFERENCES BETWEEN our earlier work that used vectors and our new approach based on color changes over time are a matter of perspective. It is the difference between going with the flow and standing still amid the current, and that change in viewpoint is what makes our newer calculations simpler to do. The idea comes from scientists who watch fluids and model how they move. There are two contrasting ways to do this: Lagrangian and Eulerian methods. Lagrangian approaches track a given portion of fluid as it travels through space, like an observer on a boat following the flow of a river. In contrast, Eulerian approaches use a fixed location in space and study the fluid passing by it, as if the observer was standing on a bridge.

Our earlier work followed a Lagrangian philosophy, acting like the observer on a boat, where pixels are tracked in the input video and then moved—as the boat moves—according to magnified vectors from point to point to point. In contrast, our new approach considers color changes only at a fixed location, similar to the observer who stays on the bridge. This local perspective applies to only small motions but makes it much simpler and robust. A computer can quickly process a video using this technique, whereas our earlier work required a lot of computer-processing time and often contained mistakes.

In 2012 we published a paper on this new method, called Eulerian video magnification. It showed how the blood flow changed a face. It also contained a variety of other examples, such as the breathing motions of an infant, which can be amplified so that parents of newborns could check an enhanced video signal to see if the baby was moving. We also took a high-speed video of a guitar where all the strings were vibrating and selected narrow bands of frequencies around given notes, such as 72 to 92 hertz for a low E string vibrating at 82 Hz. This amplified the motion of a single string, whereas the others looked like they were absolutely still.

We created a Web site where people could upload their videos and run them through this motion-magnification process. (See the videoscope at https://videoscope.qrilab.com.) People used it in ways we had not thought about, which was exciting. One person posted a video showing fetal movements in a late-term pregnancy. Another person amplified the breathing motion of her pet guinea pig. An art student made a video showing the imperceptible movements and expressions of her friends trying to stay still.

We also learned, however, that our Eulerian approach does have limits. If a given input pixel gets much darker from one frame to the next, the computer will enhance this change to an excessive degree, producing a fully black pixel, kind of a runaway amplification effect. This type of issue can cause dark or bright halos around motion areas. Input color variations from sensor noise are also a challenge because—even though we smooth them out by averaging many local pixels—the noise still gets magnified.

This result prompted us and our graduate student, Neal Wadhwa, to develop a new algorithm that preserves the benefits of simple Eulerian approaches but provides a better view when changes get more extreme.

We realized that the root of our original method's limitations is a false assumption. It acted as if the color difference between each pixel and all its neighbors—pixels to the left, the right, above, below—was the same, which unfortunately is not always true. Edges, for example, correspond with much bigger pixel differences (higher gradients) than their surrounding smooth areas. So if

SCIENTIFIC AMERICAN ONLINE Watch motion microscope videos at ScientificAmerican.com/jan2015/motion-microscope

you try and amplify all the pixels by the same amount at the same time, you get distortions that do not reveal actual movements.

Instead of amplifying by the same amount, we decided to represent each segment of an image—a local group of pixels—mathematically as a sine wave. A sine wave goes up and down, and the steep slope shows fast variation, whereas the top and bottom show a slow change. In a video image, edges mimic the fast varying part, and smooth areas look like the slow part. We can represent the change in an area of an image over time as the change in what is known as the phase of the wave. Moving from a fast varying phase to a slower varying phase helps us characterize how much movement happened between two frames of a video, and it does not create video artifacts such as halos. We reported this advance in 2013.

SMALL MOTIONS, BIG IMPLICATIONS

AFTER WORKING OUT these kinks, we found that we could process videos to see infinitesimal movements that had previously been predicted only by equations or by computer simulations. For instance, the shell surrounding the round frame of a PVC pipe is a simple object. When hit by a hammer or something similar, the shell bends and rebounds in specific patterns that oscillate at different time frequencies. Patterns that vibrate up and down quickly are tightly bunched, whereas those that shift more slowly are bigger, and they force the shell into different shapes. These patterns appear as equations in engineering textbooks, but seeing the actual deformations in the pipe was difficult because the changes are so small.

We took high-speed videos of pipes being hit. In the unprocessed video, any change in the circular shape is barely visible. Then Justin Chen, a graduate student of Oral Buyukozturk at M.I.T., working on a project with scientists at Shell International E&P, ran the video through our motion microscope, telling the computer to pull out the three lowest-frequency modes of oscillation. (This is the same principle we used to visualize a human pulse, by looking for only the pixel changes that corresponded to the heartbeat rate per minute.) Amplifying those frequencies showed the pipe cross section flexing inward and out, revealing the actual movements.

Watching a wineglass break under sound pressure—vibrating at high frequency—is another great example of how dramatic this visualization can be. We have all seen Hollywood movies when a soprano hits a high note and shatters glass. But none of us had ever seen the actual deformation of the glass because it is usually both too small in amplitude and too fast, typically around 300 to 500 Hz. We wanted to show the glass bend in and out in real time.

To do this, we used an old trick from Harold Edgerton, a pioneer of strobe-light, stop-motion photography. He showed that when a fast periodic motion is recorded with short exposures for each frame, the motion extends for several periods between frames and appears much slower than in real life. We used a regular video camera to grab short bursts of images of the glass. When we magnified the video through our motion microscope, this strobelike effect allowed us to see a glass vibrate in front of our eyes as soon as we hit it with the proper note.

The structural failure of wineglasses can disrupt a dinner party, but we hope that the motion microscope can reveal more serious effects, such as a large, potentially dangerous machine that is beginning to fail. The microscope can take small motions that may be characteristic of mechanical failure and make them visible. We showed this principle in a high-speed video of a car idling normally. As with the pipe, the raw video shows absolutely no movement of any mechanical parts. We then filtered the video to focus on engine vibrations at 22 Hz, blocking out all other frequencies. Magnifying the filtered changes by a factor of 30 revealed that different components of the engine were shaking back and forth. This was not abnormal for an engine, but it shows that the motion microscope can pick out particular bands that might be anomalous, in addition to amplifying small changes until they look big enough to see. Such videos could highlight and help diagnose malfunctioning mechanical parts in rotating or vibrating machinery.

We used a similar approach to show a giant construction crane shaking in the wind. Though appearing rigid to observers, the motion microscope shows the crane bending. There is a normal range of motion for such cranes. If the crane exceeded that range, it could spell trouble. We are exploring structure monitoring with Shell scientists Dirk Smit and Sergio Kapusta.

We can also reverse-engineer the process. By using the motion microscope to highlight tiny vibrations of objects such as the leaves of a plant, we, along with Wadhwa and Abe Davis of M.I.T. and Gautham Mysore of Adobe Research, reconstructed the type of sound that was causing them to shake. If this method was applied to, say, the concrete ramp of a bus terminal, it might identify the sources of vibrations that could weaken the structure.

The motion microscope can also be used to reveal problems in fluid flow. When the smooth flow of air or water in two adjacent layers changes into turbulent mixing, an unstable wave can form where the two layers meet. When this turbulence forms around vehicles, ranging from cars to airplanes to submarines, it has dramatic effects on how fast they move. So studying them is quite important. The waves cannot be seen in unprocessed video, but when particular motion frequencies are magnified by a factor of 40 in our video microscope, signs of wave instability pop out at the viewer.

Using the software to reveal the unseen can feel like putting on magic glasses or suddenly acquiring superhuman vision. Yet it is not magical or the dream of a comic-book creator; it is the result of basic research into video processing and mathematical representation of images. Already it has shown scientists phenomena we knew about intellectually but had never seen with our own eyes. It could, like the first optical microscopes centuries ago, help people identify threats to health and safety. Right now it makes us feel like explorers, marveling at a whole new world of phenomena that have always been around us, hidden in plain sight.

MORE TO EXPLORE

Eulerian Video Magnification for Revealing Subtle Changes in the World. Hao-Yu Wu et al. in *Proceedings of ACM SIGGRAPH 2012*, Vol. 31, No. 4, Article No. 65; July 2012. Preprint available at http://people.csail.mit.edu/mrub/vidmag

2012 International Science & Engineering Visualization Challenge. *Science*, Vol. 339, pages 518–519; February 1, 2013. www.sciencemag.org/content/339/6119/518.full

Videoscope. Quanta Research Institute. Upload your own videos to the motion microscope: https://videoscope.qrilab.com

FROM OUR ARCHIVES

Vision by Man and Machine. Tomaso Poggio; April 1984.

scientificamerican.com/magazine/sa