

Automating the Recognition of Stress and Emotion:

From Lab to Real-World Impact

It was the 1990s, and I was a new professor at MIT, trained in electrical engineering and computer science. I was very much enjoying my work on mathematical modeling and perceptual intelligence—enabling computers to perceive the varied visual and auditory streams that we humans perceive. Over vacation, I read Richard Cytowic’s *The Man Who Tasted Shapes* (MIT Press, 1993). The book described synesthesia, a kind of multimedia experience inside your brain, where people experience involuntary associations between two senses—for example, they might feel shapes in their hands while tasting foods. More commonly, people see color-letter associations that are stable across a lifetime, where an A might be red, a B green, and so on. The bizarre sensory associations should have involved the perceptual areas of the brain I was modeling computationally, in the *cortex*—or so I thought. But they didn’t. They involved areas deeper in the brain.

These deeper regions had been ignored in computational perception, because scientists figured that most intelligence was “up in the cortex.” The deep-brain regions were involved in emotion, which generally wasn’t associated with intelligence. It was easy to see why people weren’t interested in these regions. However, although I wasn’t interested in emotion and had no desire to become associated with it—an association that I figured would undo all of my hard work in building a respectable reputation—I had to admit that there were more pathways going from these deep “emotion” regions of the brain to the regions I was modeling than vice-versa. So, I decided to (quietly) learn more.

The rest of the story of my personal adventures, which led to a book, *Affective Computing* (MIT Press, 1997), is told in the opening article

of what became the first international journal of the new field, *IEEE Transactions on Affective Computing*.¹ Here, I want to describe newer adventures, where our work moved from research best characterized as, “Is this even possible?” to “Aargh, I am getting too many emails asking for the technology we created, and the requests are for great causes; how do I take care of these and get back to research?”

Measuring the Physiology of Stress

Since the late 1800s, scientists have debated what emotion is—in particular, whether emotions are cognitive constructs, or whether they have a unique physiological pattern associated with them. A common belief was that there was just “general arousal” in the body that provided a feeling, and all the things that differentiate emotion—for example, whether it is positive or negative—were simply cognitions. Along with Jennifer Healey, my first PhD student willing to work on emotion, I set out to find reliable ways to elicit a set of emotions and see if, by measuring multiple modalities (such as muscle tension, respiration, skin conductance, and heart-rate variability), we could identify any patterns in emotion that could be recognized reliably within a person.

Understanding the Limits of Lab Data

In our first person-dependent, long-term effort with lab-based measurements, we successfully collected 30 days of data and automated the recognition of eight emotions (including neutral) with 81 percent accuracy (see Figure 1).² This was a breakthrough, showing that some kind of automated emotion recognition was possible. We also showed that it wasn’t just arousal being recognized.

Rosalind W. Picard
MIT Media Lab



Figure 1. Jennifer Healey demonstrates the first affective measurement system at MIT, used to collect data to automatically identify eight emotions from a person over 30 days.




Figure 2. Our first automated system to automatically measure driver stress.³

I spoke with our experimental subject to ask about her experiences. She said the anger she felt in the lab was nothing compared with what she felt when she left the lab. Although we could elicit some aspects of emotion in the lab, we needed to study it in the real world, where “what matters to a person” happens. My student, Steve Mann, had been building wearable computers and cameras that modified his perceptual experience. Healey and I decided to build the first “wearable affective computer” that could measure your affective state in real

life, modify your computer’s response, and hopefully improve your affective experience.

Studying Real-World Data

Stress is a key component of affective experience—and a highly relevant emotion at MIT and in Boston (in particular, in Boston driving). The physiology of stress is complex. There is no one gold standard for “truth,” so we set out to measure it in multiple ways, with video, multiple physiology channels, self-reported feelings, and observer ratings of Boston drivers and of the complexity of their driving situations.

Although the driving situation is nicely constrained because the drivers are seated behind the wheel of a car, they still move a lot, making it a challenge to get clean data. Healey built our sensor system and structured a task with a series of resting and driving segments, some relatively relaxed and some enormously stressful (see Figure 2). We only had one driver get in an accident, and everyone was fine. We eventually obtained the world’s first set of rich real-world multimodal driver stress data over 24 trips across Boston and surrounding towns. Importantly, we also collected contextual data, including bumpy encounters with deep potholes, and near encounters with pedestrians who strolled out in front of the car, nowhere near a crosswalk.

What did we learn about stress? Boston driver stress had nothing to do with the (reasonable) speeds people drove (city or highway) in our study, and a lot to do with uncertainty within a situation and the complexity of the context.³ Combining multiple ground truths, we were able to examine low-, medium-, and high-stress conditions. Overall, our data showed that a single modality—skin conductance—measured from the electrodermal activity (EDA) gave the highest correlation with our multiple measures of stress. EDA reflects activity in the sympathetic nervous system branch of the autonomic nervous system—a kind of “autonomic stress.”⁴ Heart rate and heart-rate variability, which also capture different kinds of autonomic stress, were also sometimes helpful.

With this knowledge, in 2013 we moved on to new work, as I started learning more about the enormous stress and anxiety experienced by many people with autism.

Building a Wearable EDA Sensor

People with heightened sensitivity to sounds, fluorescent lights, fragrances in everyday

products, eye contact, and other environmental stimuli, can easily become overloaded and can “shut down.” Many people with autism suffer from such overload, which can also lead to behaviors that are injurious to the self or others. So, in an effort to better understand those with autism, I decided to build a wearable EDA sensor that could continuously measure and wirelessly communicate their autonomic stress.

Together with Rich Fletcher and a team of students, I built sensors into sweatbands worn on wrists and ankles. We were amazed to see the measured data climb before some meltdowns and during tasks that increased cognitive or physical exertion. We watched the skin conductance level fall, like a slide on a playground, with repetitive movements like swinging or rocking. Suddenly, a nonspeaking person had a way to show what might be stressing her out, or bringing on calm.

Is This Possible?

Shortly before the winter break, an undergraduate asked me if he could borrow one of our sensors to see what was causing stress for his autistic little brother, who is nonspeaking. I said, “Sure, take two,” since back then the wires would often break. He put one on each wrist of his little brother and watched the data stream wirelessly in real time.

Later, as I sat in my office reviewing the boy’s data, I thought, “This day looks pretty typical,” and “normal variation here (yawn).” Both wrists transmitted signals that went up with mild autonomic stressors, and down with relaxation. Usually, the two sides of the body responded with similar signals, and everything looked normal. Then I clicked to see the next day’s data, and my jaw dropped. One wrist showed a peak that was greater than 10 times the typical stress response. His other wrist showed no response at all. My first thought was that one of the sensors was broken. After all, how can you have stress on just one side of your body? And that large?

I am an electrical engineer, so I began debugging. Nothing sensor-related explained what I saw. In fact, the data right before and after this weird episode showed normal behavior on both sides, with a clear “sleep signature.” I have probably looked at more electrodermal data than anybody on the planet, and I could not think of anything explaining this; I was perplexed.

The next day I did something I’d never done: I called a student at home on his vacation.

“Hi, how was your Christmas? How is your little brother? Hey, any idea what happened to him at [exact date and time of the stress response].” The student didn’t know but said he’d check his diary. (“Diary? An MIT student keeps a diary?”) I held my hands together in prayer, thinking the odds were nil he’d have written down this moment of his multi-week vacation. He returned, confirmed the time and date with me, and told me that was 20 minutes before his little brother had a grand mal seizure.

A giant signal on the wrist before a seizure?

Conducting a Real-World Study

I called the chief of neurosurgery, Joseph Madsen, at Children’s Hospital Boston (CHB). “Hi, Dr. Madsen. My name is Rosalind Picard... Is it possible there could be a huge sympathetic nervous system surge 20 minutes before a seizure?” I didn’t want to tell him it was just on one side of the body. After all, we were measuring a component of emotion. Emotion on only one side? I didn’t want him to hang up on me.

Madsen was very nice, “Probably not possible 20 minutes before the seizure starts in the brain. But it might happen before the outward clinical signs.” Then he paused. “I have seen patients have their hair stand on end on only one arm before a seizure. Or have goosebumps on only one side.” On *one* side? I told him about the asymmetry, and he got very interested.

Then, after getting approval from MIT and CHB, we made more wristbands and ran a study that simultaneously measured EEG, ECG, video, and also EDA. Ming-Zher Poh, a doctoral student at MIT, designed and built better sensors for logging quality data 24/7, conducting this risky seizure research for his PhD.

What did we find? The doctors labeled the patient’s videos and EEGs for seizures while blinded to our data. We found that 100 percent of the most severe seizures, called “generalized tonic clonic (GTC)” or “grand mal,” had significant EDA responses. Also, 86 percent of the “complex partial seizures,” which don’t have convulsions but cause the patient to lose consciousness, showed EDA surges more than two standard deviations above the average pre-seizure period. In most cases, the seizures were generalized to both sides of the brain, and the wristband responses were on both sides of the body. Unfortunately, the EDA responses were not usually in advance of the seizures when we had precise timing; they usually started on the wrist seconds after the seizure started in the



Figure 3. *The Empatica Embrace wristband measures multiple physiological signals and can run apps to detect and communicate autonomic stress and movement patterns, potentially saving lives.*

brain. But there was another very big—even more important—surprise.

Many of our patients had a period of time after the seizure ended when all their brain waves (measured on the scalp via EEG) went flat. This is called PGES (Post-ictal generalized EEG suppression). The EEG showed that the seizure “ended,” but the brain waves, instead of going back to normal, looked “dead.” Fortunately, nobody died. However, these prolonged periods of flat brain wave activity after a seizure have been observed in all monitored cases of sudden unexpected death from epilepsy (SUDEP).

SUDEP takes more lives in the US than house fires or sudden infant death syndrome. It is the number one cause of death in epilepsy, and when it happens, it usually occurs many minutes after the seizure appears to have ended, when a person might be left alone to sleep.^{5,6}

We found that a signal measured by the wristband was highly correlated with how long the brain waves were suppressed after the seizure. In other words, the bigger the signal on the wrist, the longer the brain waves went flat after the seizure had supposedly ended.⁷

Usually you need to wear an EEG to detect brain wave suppression. Wearing an EEG is inconvenient, uncomfortable, and not stylish (except perhaps at MIT parties). We had found a useful correlate of the EEG suppression in a comfortable wristband.

Detecting Convulsive Seizures

We also learned some other amazing things about the human brain. A key part of the brain involved in emotion is the amygdala: we have two—one on the right and one on the left. When either amygdala gets electrically stimu-

lated (this requires invasive procedures), then it elicits a large skin conductance response on the palm on the same side of the body.⁸ When either amygdala gets stimulated with 15–20 volts in a sustained way, it causes the person to stop breathing.⁹ Moreover, the person *can* breathe; he or she just doesn’t. But, if you ask a question, prompting the person to try to talk, then he or she starts breathing again.

Recent findings showed that 100 percent of observed SUDEPS began when the patient stopped breathing.⁶ One possible explanation for this is that the seizure spread to amygdala, activating it in a way that turned off the person’s breathing. If this happens, then patients might need somebody to come near and touch or talk to them to help them start breathing again.

Although we had set out to measure emotional stress on the wrist, the patterns picked up by our wristband were indicating atypical brain activity, deep in the brain, even though the EEG showed no brain activity on the scalp. Poh was able to use the data to build an accurate automated detector of convulsive seizures.¹⁰ These findings also led to studies of patients in coma after cardiac arrest, where EDA was shown to help determine who survives.¹¹

Today, a company that I co-founded, Empatica Inc., has commercialized these capabilities into a wristband that can measure the clinical quality of the data we need to do our research, run on-board machine learning and pattern analysis, and issue alerts to caregivers (see Figure 3). I got an email recently from one beta user, who received an alert and found her daughter face-down in bed after a short seizure not breathing. After she turned her daughter over, she started breathing again. The mom emailed me enthusiastically with the news, sending pictures of her little girl now “pink” and happily playing.

Reflecting on these events, there are many surprises: Who would have expected that our efforts to develop machine perception would lead to a wearable that detects signals related to deep brain activation, and issues potentially life-saving alerts? The life of an engineer can be full of adventure when following a few guidelines: First, get good data. Second, keep trying to understand the data, especially when it looks bizarre or wrong. Third, be fearless, even if it means picking up the phone to call the chief of neurosurgery or a student at home on vacation. **MM**

References

1. R.W. Picard, "Affective Computing: From Laughter to IEEE," *IEEE Trans. Affective Computing*, vol. 1, no. 1, 2010, pp. 11–17.
2. R.W. Picard, E. Vyzas, and J. Healey, "Toward Machine Emotional Intelligence: Analysis of Affective Physiological State," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23, no. 10, 2001, pp. 1175–1191.
3. J.A. Healey and R.W. Picard, "Detecting Stress During Real-World Driving Tasks Using Physiological Sensors," *IEEE Trans. Intelligent Transportation Systems*, vol. 6, no. 2, 2005, pp. 156–166.
4. W. Boucsein, *Electrodermal Activity*, Springer, 2012.
5. M.J. England et al., "Epilepsy Across the Spectrum: Promoting Health and Understanding: A Summary of the Institute of Medicine Report," *Epilepsy & Behavior*, vol. 25, no. 2, 2012, pp. 266–276.
6. P. Ryvlin et al., "Incidence and Mechanisms of Cardiorespiratory Arrests in Epilepsy Monitoring Units (MORTEMUS): A Retrospective Study," *The Lancet Neurology*, vol. 12, no. 10, 2013, pp. 966–977.
7. M.-Z. Poh et al., "Autonomic Changes with Seizures Correlate with Postictal EEG Suppression," *Neurology*, vol. 78, no. 23, 2012, pp. 1868–1876.
8. C.A. Mangina and J.H. Beuzeron-Mangina, "Direct Electrical Stimulation of Specific Human Brain Structures and Bilateral Electrodermal Activity," *Int'l J. Psychophysiology*, vol. 22, no. 1, 1996, pp. 1–8.
9. B.J. Dlouhy et al., "Breathing Inhibited when Seizures Spread to the Amygdala and upon Amygdala Stimulation," *J. Neuroscience*, vol. 35, no. 28, 2015, pp. 10281–10289.
10. M.-Z. Poh et al., "Convulsive Seizure Detection Using a Wrist-Worn Electrodermal Activity and Accelerometry Biosensor," *Epilepsia*, vol. 53, no. 5, 2012, pp. e93–e97.
11. V. Alvarez et al., "Continuous Electrodermal Activity as a Potential Novel Neurophysiological Biomarker of Prognosis after Cardiac Arrest—A Pilot Study," *Resuscitation*, vol. 93, Aug. 2015, pp. 128–135.

Rosalind W. Picard is the director of affective computing research and a professor at the MIT Media Lab. She is also co-founder, chairman, and chief scientist of Empatica. Contact her at picard@media.mit.edu or at rp@empatica.com.



COMPUTER ENTREPRENEUR AWARD

In 1982, on the occasion of its thirtieth anniversary, the IEEE Computer Society established the Computer Entrepreneur Award to recognize and honor the technical managers and entrepreneurial leaders who are responsible for the growth of some segment of the computer industry. The efforts must have taken place over fifteen years earlier, and the industry effects must be generally and openly visible.

All members of the profession are invited to nominate a colleague who they consider most eligible to be considered for this award. Awarded to individuals whose entrepreneurial leadership is responsible for the growth of some segment of the computer industry.

**DEADLINE FOR 2017 AWARD NOMINATIONS
DUE: 15 OCTOBER 2016**

AWARD SITE: <https://www.computer.org/web/awards/entrepreneur>
www.computer.org/awards

