SRMS College of Engineering & Technology, Bareilly

Vol. 2
Issue-4
Sep. 2016

E-NEWS LETTER

Computer Applications

CONTENTS

LATEST UPDATES WORLD

Automated Screening for Childhood Communication Disorders

For children with speech and language disorders, early-childhood intervention can make a great difference in their later academic and social success. But many such children -- one study estimates 60 percent -- go undiagnosed until kindergarten or even later.

A new computer system can automatically screen young children for speech and language disorders and, potentially, even provide specific diagnoses.

Researchers at the Computer Science and Artificial Intelligence Laboratory at MIT and Massachusetts General Hospital's Institute of Health Professions hope to change that, with a computer system that can automatically screen young children for speech and language disorders and, potentially, even provide specific diagnoses.

This week, at the Interspeech conference on speech processing, the researchers reported on an initial set of experiments with their system, which yielded promising results. "We're nowhere near finished with this work," says John Guttag, the Dugald C. Jackson Professor in Electrical Engineering and senior author on the new paper. "This is sort of a preliminary study. But I think it's a pretty convincing feasibility study."

The system analyzes audio recordings of children's performances on a standardized storytelling test, in which they are presented with a series of images and an accompanying narrative, and then asked to retell the story in their own words.

"The really exciting idea here is to be able to do screening in a fully automated way using very simplistic tools," Guttag says. "You could imagine the storytelling task being totally done with a tablet or a phone. I think this opens up the possibility of low-cost screening for large numbers of children, and I think that if we could do that, it would be a great boon to society."

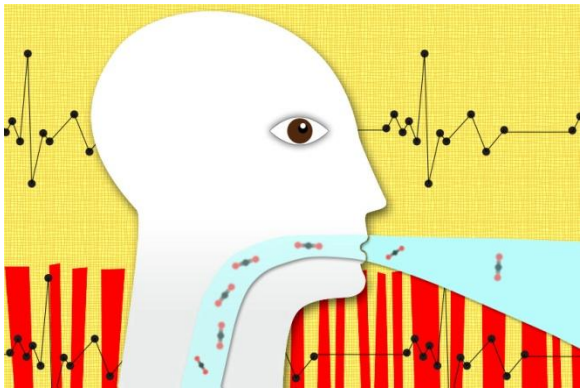

Subtle signals

The researchers evaluated the system's performance using a standard measure called area under the curve, which describes the tradeoff between exhaustively identifying members of a population who have a particular disorder, and limiting false positives. (Modifying the system to limit false positives generally results in limiting true positives, too.) In the medical literature, a diagnostic test with an area under the curve of about 0.7 is generally considered accurate enough to be useful; on three distinct clinically useful tasks, the researchers' system ranged between 0.74 and 0.86.

To build the new system, Guttag and Jen Gong, a graduate student in electrical engineering and computer science and first author on the new paper, used machine learning, in which a computer searches large sets of training data for patterns that correspond to particular classifications -- in this case, diagnoses of speech and language disorders.

The training data had been amassed by Jordan Green and Tiffany Hogan, researchers at the MGH Institute of Health Professions, who were interested in

developing more objective methods for assessing results of the storytelling test. "Better diagnostic tools are needed to help clinicians with their assessments," says Green, himself a speech-language pathologist. "Assessing children's speech is particularly challenging because of high levels of variation even among typically developing children. You get five clinicians in the room and you might get five different answers."

Unlike speech impediments that result from anatomical characteristics such as cleft palates, speech disorders and language disorders both have neurological bases. But, Green explains, they affect different neural pathways: Speech disorders affect the motor pathways, while language disorders affect the cognitive and linguistic pathways.

Telltale pauses

Green and Hogan had hypothesized that pauses in children's speech, as they struggled to either find a word or string together the motor controls required to produce it, were a source of useful diagnostic data. So that's what Gong and Guttag concentrated on. They identified a set of 13 acoustic features of children's speech that their machine-learning system could search, seeking patterns that correlated with particular diagnoses. These were things like the number of short and long pauses, the average length of the pauses, the variability of their length, and similar statistics on uninterrupted utterances.

The children whose performances on the storytelling task were recorded in the data set had been classified as typically developing, as suffering from a language impairment, or as suffering from a speech impairment. The machine-learning system was trained on three different tasks: identifying any impairment, whether speech or language; identifying language impairments; and identifying speech impairments.

One obstacle the researchers had to confront was that the age range of the typically developing children in the data set was narrower than that of the children with impairments: Because impairments are comparatively rare, the researchers had to venture outside their target age range to collect data.

Gong addressed this problem using a statistical technique called residual analysis. First, she identified correlations between subjects' age and gender and the acoustic features of their speech; then, for every feature, she corrected for those correlations before feeding the data to the machine-learning algorithm.

"The need for reliable measures for screening young children at high risk for speech and language disorders has been discussed by early educators for decades," says Thomas Campbell, a professor of behavioral and brain sciences at the University of Texas at Dallas and executive director of the university's Callier Center for Communication Disorders. "The researchers' automated approach to screening provides an exciting technological advancement that could prove to be a breakthrough in speech

and language screening of thousands of young children across the United States."

Giving Voice to Emotion

It's tough to imagine anything more frustrating than interacting with a call center. Generally, people don't reach out to call centers when they're happy—they're usually trying to get help with a problem or gearing up to do battle over a billing error. Add in an automatic phone tree, and you have a recipe for annoyance. But what if that robotic voice offering you a smorgasbord of numbered choices could tell that you were frustrated and then funnel you to an actual human being? This type of voice analysis technology exists, and it's just one example of the many ways that computers can use your voice to extract information about your mental and emotional state—including information you may not think of as being accessible through your voice alone.

Detecting Emotions

Ten years ago, when Izhak Shafran was a researcher at AT&T Research Labs, he and his team wanted to know whether different characteristics of a customer's voice—so-called voice signatures—could tell them information about that customer, such as gender, age, dialect, and emotion. Using actual speech collected from AT&T's "How May I Help You" customer call system, Shafran and his colleagues were able to train an algorithm to detect these characteristics at levels high above chance.

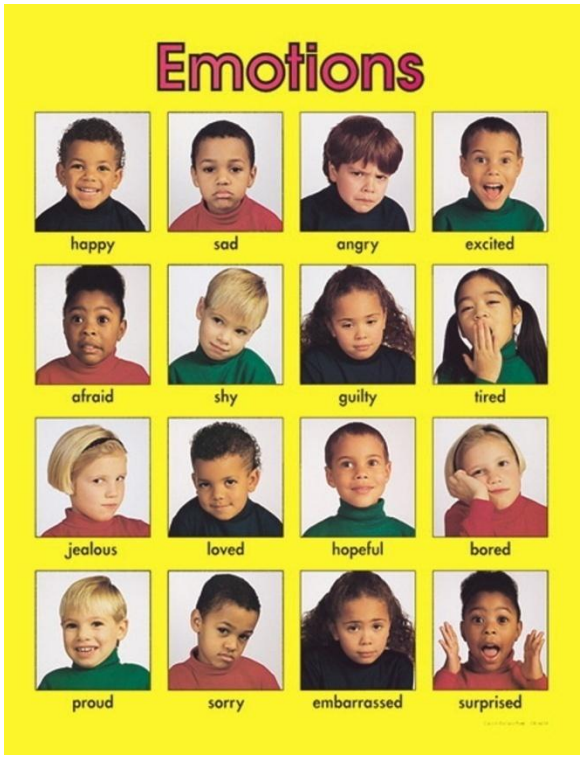

Using only information about pitch and something called the *Mel frequency cepstral coefficient*, a vocal feature that is commonly used in voice recognition tools, the algorithm was able to correctly identify a caller's gender with 95% accuracy, approximate age with 70% accuracy, dialect with 45% accuracy, and emotion with 68% accuracy. "We could easily detect things like frustration," says Shafran, who is now a speech researcher at Google.

Fast forward to the present when a Boston-based company called Cogito uses voice analysis algorithms to help actual humans—customer service agents—gain real-time insight into how their conversations are going by analyzing several different features, such as the degree of a customer's pitch variation, which can indicate boredom or anger. Cogito has an interface that allows

customer service agents to look at their own voice features, too, so that they are able to "dynamically adjust their style to align with the customer's preference."

Several academic groups are also working on applications that can extract emotional information from speech samples. One example is EmoVoice, a "comprehensive framework for real-time recognition of emotions from acoustic properties of speech," developed by Elisabeth André, Ph.D., and her team at the University of Augsburg in Germany. EmoVoice performs the same three steps used by most tools designed to recognize emotions from speech: audio segmentat ion (breaking down speech samples into pieces that can be analyzed), feature extraction (finding which acoustic characteristics best describe emotions), and classification (using machine learning and statistical modeling techniques to train algorithms to detect which extracted features are associated with which emotions).

The first step, audio segmentation, can be tricky. Words are unlikely to be long enough to be useful, but if you analyze too large a chunk of speech, features can be washed out. So Emo- Voice divides audio samples into chunks akin to phrases—not too short, but not too long. The second step, feature extraction, requires taking measurements and determining the acoustic properties that can best characterize emotions. Examples of such properties are pitch and voice intensity (or loudness). EmoVoice can extract 1,302 features, although only between 50 and 200 are used when analyzing any given sample. The third step, classification, uses computer algorithms to sort the extracted features into groups representing different emotions. For example, a monotonous audio sample that has minimal pitch variation could be grouped as an indicator of sadness or boredom.

EmoVoice, which is available to everyone, has already been integrated into several applications. These include a humanoid robot named Barthoc, who can express joyful and fearful facial expressions when listening to a fairy tale; a virtual agent named Greta, who can mirror the emotions of a speaker with her facial expressions and deliver appropriate verbal feedback; and art installations with emotional kaleidoscopes and a tree that grows and changes color and shape based on the emotions extracted from people's voices.

STUDENTS CORNER

10 Breakthrough Technologies 2016

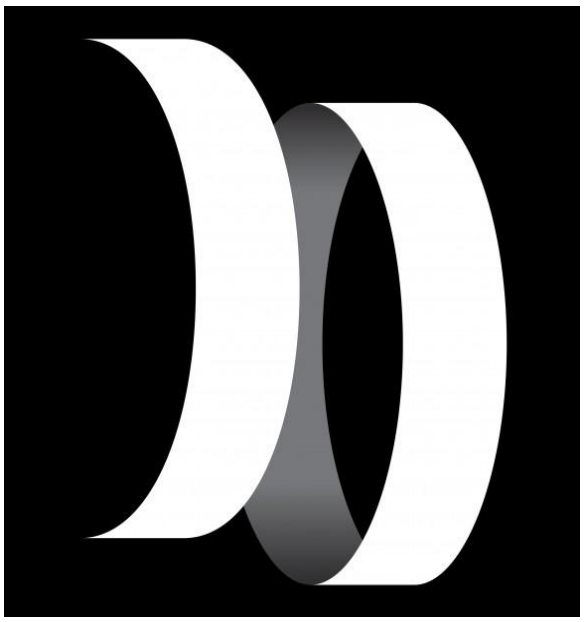

Which of today's emerging technologies have a chance at solving a big problem and opening up new opportunities? Here are our picks. The 10 on this list all had an impressive milestone in the past year or are

on the verge of one. These are technologies you need to know about right now.

Immune Engineering

Genetically engineered immune cells are saving the lives of cancer patients. That may be just the start.

The doctors looking at Layla Richards saw a little girl with leukemia bubbling in her veins. She'd had bags and bags of chemotherapy and a bone marrow transplant. But the cancer still thrived. By last June, the 12-month-old was desperately ill. Her parents begged—wasn't there anything?

There was. In a freezer at her hospital—Great Ormond Street, in London—sat a vial of white blood cells. The cells had been genetically altered to hunt and destroy leukemia, but the hospital hadn't yet sought permission to test them. They were the most extensively engineered cells ever proposed as a therapy, with a total of four genetic changes, two of them introduced by the new technique of genome editing.

Precise Gene Editing in Plants

CRISPR offers an easy, exact way to alter genes to create traits such as disease resistance and drought tolerance.

A new gene-editing method is providing a precise way to modify crops in hopes of making them yield more food and resist drought and disease more effectively. Research in the past year has shown that the resulting plants have no traces of foreign DNA, making it possible that they will not fall under existing regulations governing genetically modified organisms and will sidestep many of the consumer concerns over these GMOs.

Conversational Interfaces

Powerful speech technology from China's leading Internet Company makes it much easier to use a smartphone.

Stroll through Sanlitun, a bustling neighborhood in Beijing filled with tourists, karaoke bars, and luxury shops, and you'll see plenty of people using the latest smartphones from Apple, Samsung, or Xiaomi. Look closely, however, and you might notice some of them ignoring the touch screens on these devices in favor of something much more efficient and intuitive: their voice.

A growing number of China's 691 million Smartphone users now regularly dispense with swipes, taps, and tiny keyboards when looking things up on the country's most popular search engine, Baidu. China is an ideal place for voice interfaces to take off, because Chinese characters were hardly designed with tiny touch screens in mind. But people everywhere should benefit as Baidu advances speech technology and makes voice interfaces more practical and useful. That could make it easier for anyone to communicate with the machines around us.

"I see speech approaching a point where it could become so reliable that you can just use it and not even think about it," says Andrew Ng, Baidu's chief scientist and an associate professor at Stanford University. "The best technology is often invisible, and as speech recognition becomes more reliable, I hope it will disappear into the background."

Reusable Rockets

Rockets typically are destroyed on their maiden voyage. But now they can make an upright landing and be refueled for another trip, setting the stage for a new era in spaceflight.

Thousands of rockets have flown into space, but not until 2015 did one return like this: it came down upright on a landing pad, steadily firing to control its descent, almost as if a movie of its launch were being played backward. If this can be done regularly and rockets can be refueled over and over, spaceflight could become a hundred times cheaper.

Robots That Teach Each Other

What if robots could figure out more things on their own and share that knowledge among themselves?

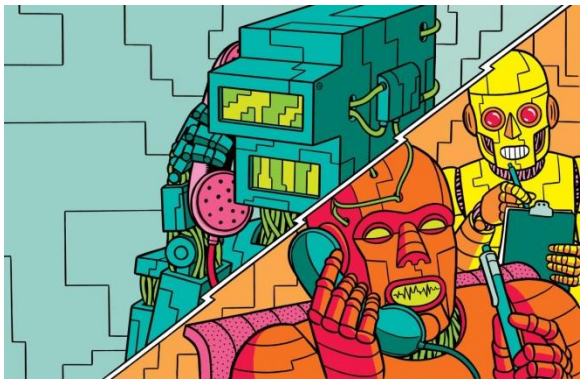

Many of the jobs humans would like robots to perform, such as packing items in warehouses, assisting bedridden patients, or aiding soldiers on the front lines, aren't yet possible because robots still don't recognize and easily handle common objects. People generally have no trouble folding socks or picking up water glasses, because we've gone through "a big data collection process" called childhood, says Stefanie Tellex, a computer science professor at Brown University. For robots to do the same types of routine tasks, they also need access to reams of data on how to grasp and manipulate objects. Where does that data come from? Typically it has come from painstaking programming. But ideally, robots could get some information from each other.

DNA App Store

An online store for information about your genes will make it cheap and easy to learn more about your health risks and predispositions.

While driving and listening to National Public Radio one day, Justin Kao heard about the discovery of a "sweet tooth gene" that makes you more likely to crave sweets. "Oh my God," thought Kao, who has always loved cookies. "I would pay $5 to know if I had that."

Kao is hoping that millions of other people will be just as eager to spend a few bucks for tidbits revealed in their DNA. He is a cofounder of Helix, a San Francisco–based company that last summer secured more than $100 million in a quest to create the first "app store" for genetic information.

SolarCity's Gigafactory

A $750 million solar facility in Buffalo will produce a gigawatt of high-efficiency solar panels per year and make the technology far more attractive to homeowners.

In an industrial park near the shore of Lake Erie, hard by the Buffalo River, the future of the solar power industry is under construction. SolarCity's sprawling Buffalo factory, built and paid for by the state of New York, is nearing completion and will soon begin producing some of the most efficient solar panels available commercially. Capable of making 10,000 solar panels a day, or one gigawatt of solar capacity a year, it will be the largest solar manufacturing plant in North America and one of the biggest in the world.

Tesla Autopilot

The electric-vehicle maker sent its cars a software update that suddenly made autonomous driving a reality.

In October 2014, Elon Musk's electric-car company began rolling out sedans with a dozen ultrasonic sensors discreetly placed around both bumpers and sides. For an additional $4,250, Tesla customers could purchase a "technology package" that used the sensors, as well as a camera, a front radar, and digitally controlled brakes, to help avoid collisions—essentially allowing the car to take over and stop before crashing. But mostly, the hardware sat there, waiting, waiting, and gathering reams of data. A year later, last October 14, the company sent a

software update to the 60,000 sensor-laden cars it had sold in that time. The software update was officially named Tesla Version 7.0, but its nickname—Autopilot—was what stuck.

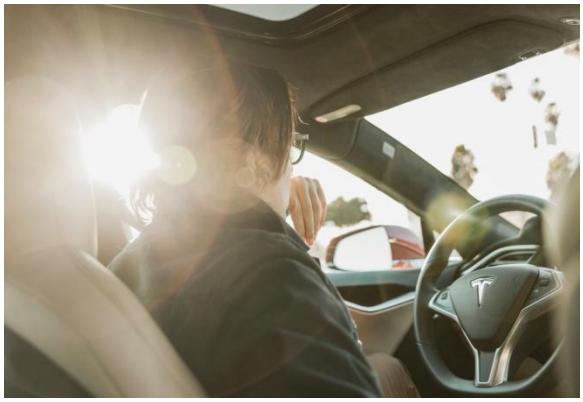

Slack

A service built for the era of mobile phones and short text messages is changing the workplace.

Deepak Gupta
MCA-2015

FACULTY ARENA

Power from the Air

[Internet devices powered by Wi-Fi and other telecommunications signals will make small computers and sensors more pervasive.]

Even the smallest Internet-connected devices typically need a battery or power cord. Not for much longer. Technology that lets gadgets work and communicate using only energy harvested from nearby TV, radio, cell-phone, or Wi-Fi signals is headed toward commercialization. The University of Washington researchers who developed the technique have demonstrated Internet-connected temperature and motion sensors, and even a camera, powered that way.

Transferring power wirelessly is not a new trick. But getting a device without a conventional power source to communicate is harder, because generating radio signals is very power-intensive and the airwaves harvested from radio, TV, and other telecommunication technologies hold little energy.

Shyamnath Gollakota and his colleague Joshua Smith have proved that weak radio signals can indeed provide all an Internet gadget needs, using a principle called backscattering. Instead of generating original signals, one of their devices selectively reflects incoming radio waves to construct a new signal—a bit like an injured hiker sending an SOS message using the sun and a mirror. A gadget using the technique absorbs some energy from the signal it is modifying to power its own circuits.

"We can get communication for free," says Gollakota. RFID chips for the contactless smart cards used in mass transit also rely on backscattering, but they require specialized reader devices and can communicate only within a few inches because the reflected signals are weak and the reader itself presents interference.

One version of the University of Washington technology, dubbed passive Wi-Fi, is being commercialized through a spin-off company, Jeeva Wireless. It lets battery-free gadgets connect with conventional devices such as computers and smartphones by backscattering Wi-Fi signals. In tests, prototype passive Wi-Fi devices have beamed data as far as 100 feet and made connections through walls. Doing that requires altering the software of a Wi-Fi access point to generate an extra signal for passive Wi-Fi devices to use, very slightly increasing its power consumption.

Smith says that passive Wi-Fi consumes just 1/10,000th as much power as existing Wi-Fi chipsets. It uses a thousandth as much power as the Bluetooth LE and ZigBee communications standards used by some small connected devices and has a longer range. A device using passive Wi-Fi to communicate—for example, a security camera—could power its other circuits using energy harvested from the Wi-Fi signals it is backscattering, or by feeding on other signals such as TV and radio broadcasts.

The researchers believe that tiny passive Wi-Fi devices could be extremely cheap to make, perhaps less than a dollar. In tomorrow's smart home, security cameras, temperature sensors, and smoke alarms should never need to have their batteries changed.

Pushpendra Kumar Singh
Asst. Professor

PARTICIPATION IN RESEARCH & DEVELOPMENT ACTIVITIES

A Research paper of **Mr. Ajeet Kumar**, entitled "**The Impact of Agile based Software Engineering in Interactive Art Installation**" is accepted and presented at **International Conference on Advances in Computing, Communication & Automation (ICACCA 2016)** which is Technically Co-Sponsored By **IEEE** held on 30 September and 01 October 2016.

PARTICIPATION IN RESEARCH & DEVELOPMENT ACTIVITIES

A Research paper of **Mr. Anuj Kumar**, entitled "**Dynamic Load Balancing in Heterogeneous Hierarchical Computational Grids Using Fuzzy Logic (LBHHGF)**" is accepted at **International Conference on Advance Computing and Software Engineering (ICACSE 2016)** which is Technically Co-Sponsored By **IEEE** & will be held on 14 -15 October 2016.

PARTICIPATION IN RESEARCH & DEVELOPMENT ACTIVITIES

A Research paper of **Mr. Ajeet Kumar**, entitled "**A Study on Developing Digital Art Patterns using High Level Programming Language**" is accepted at **International Conference on Advance Computing and Software Engineering (ICACSE 2016)** which is Technically Co-Sponsored By **IEEE** & will be held on 14 -15 October 2016.