BY ALEXANDER GELFAND

MOBILE HEALTH: BD2K CENTERS HARNESS SENSOR DATA

Having already revolutionized fields ranging from communications to finance, mobile technology and data science are now poised to do the same for healthcare.

That, at least, is the promise of the burgeoning mobile health (mHealth) movement. Thanks to the proliferation of wearable biosensors capable of recording everything from physical activity to blood oxygen levels—and the increasingly sophisticated algorithms used to sift through the mounting pile of data—researchers are finding novel ways of diagnosing illnesses, predicting disease risk, and promoting healthier lifestyles.

Moreover, two NIH Big Data to Knowledge Centers of Excellence—the **Mobilize Center** and the Center **Mobile Sensor Data-to-Knowledge (MD2K)**—are paving the way for the entire mHealth community.

"The research methodologies, algorithms and devices these centers are developing—not to mention the training opportunities they provide—are creating a foundation that will make it easier for others to produce robust mobile health research," says **Scott Delp, PhD**, professor of bioengineering at Stanford University and principal investigator of the Mobilize Center.

Disease Detection with Smart Devices

Some Mobilize Center researchers are leveraging consumer products that are already used by large numbers of people. **Jessilyn Dunn, PhD**, a postdoctoral fellow at Stanford University, recently evaluated the possibility of using commercially available wearables to gather and analyze health-related information in ways that aren't normally done in the clinic.

As reported in an article published in January 2017 in *PLoS Biology*, Dunn and her colleagues, including **Michael Snyder, PhD**, professor of genetics at Stanford and director of the Center for Genomics and Personalized Medicine, performed several different experiments using a variety of wearables. They found that two commonly used tools provided most of the information they needed: a smartwatch capable of detecting heart rate, skin temperature, and activity; and a smartphone capable of reporting activity and location.

When combined with the occasional use of a wearable oxygen sensor, these devices collected much of the same information that would ordinarily be recorded in a doctor's office once every year or so. In this case, however, the data were gathered regularly—often continuously—over a lengthy period: one study participant was monitored for two years, while an additional 43 participants were monitored for an average of 11 months.

Reconciling the different data formats used by competing companies wasn't easy, especially when manufacturers periodically changed the way they packaged the output from their products. As a result, Dunn says, the team spent a lot of their time cleaning the data and "making sure that everything was kosher from one iteration to the next" to ensure they were "comparing apples to apples."

Participants also underwent blood testing on a regular basis. This allowed the researchers to unearth several interesting findings that were lurking in the sensor data.

For example, the researchers retrospectively detected the onset of a viral infection—in this case Lyme disease—in one subject based solely on elevated skin temperature and unusual heart-rate patterns. This prediction was confirmed by the presence of Lyme bacteria antibodies in his blood. Delving more deeply into the subject's data, the researchers identified several other periods of illness during which similar abnormalities in temperature and heart rate were accompanied by the presence of an inflammatory biomarker known as high-sensitivity C-reactive protein in his blood.

Based on his data, the researchers developed an algorithm, called Change-of-Heart, that identified instances of illness amongst several other participants before they reported symptoms, based solely on abnormalities in their heart rates.

"The research methodologies, algorithms and devices these centers are developing—not to mention the training opportunities they provide—are creating a foundation that will make it easier for others to produce robust mobile health research," says Delp.

Eventually, that kind of predictive capability could allow sensor-based systems to warn people of an impending illness even before they feel sick—enabling an algorithm to tell you to "run to your local pharmacy and pick up some cold medicine, because tomorrow you're going to wake up with a cold," Dunn says.

In a similar vein, Dunn and her

RELEVANT NIH INSTITUTES:

NIMH, NINDS, NHLBI, NIA, NIBIB, NIDA, NHGRI, NICHD, and NIAAA

colleagues successfully identified sensor-based predictors of insulin resistance, a risk factor for type II diabetes, which they confirmed by testing steady-state plasma glucose (SSPG) levels among a subset of study participants.

The researchers started with clinically measured body mass index (BMI), and added sensor-based reports of both physical activity and heart rate—in particular, differences between day and nighttime heart-rate patterns that they found to be associated with diabetes. While the researchers found that they could best predict insulin resistance if they used all three parameters in combination, variation in heart rate proved to be the strongest biomarker of the lot, and was an effective predictor even in the absence of the other two.

Given their usefulness, Dunn hopes that as wearables become cheaper, they will help expand healthcare access to low-income groups and people in remote rural communities, many of whom cannot easily see a living, breathing doctor. "This is really a fantastic public health opportunity," she says.

In the meantime, the rich dataset she and Snyder created is available (at http://hmpdacc.org/data/wearable/stanford.tar) for others to explore.

Boosting Health with Games and Social Networks

Tim Althoff, MS, a doctoral candidate in computer science who is also affiliated with the Mobilize Center, is trying to promote population-scale health benefits as well. But his tool of choice is the online social network. And rather than dealing with dozens of study participants, he's working with thousands, even millions, of them.

In a paper published in the *Journal of Medical Internet Research* in 2016, Althoff and his colleagues demonstrated that users of the Pokémon Go augmented-reality game—a group whom Althoff describes as being "way less active than the typical health-app user"—did in fact become more physically active because of their involvement with the app.

Althoff and his co-authors had access to data from nearly 32,000 users of the Microsoft Band wearable fitness tracker. By algorithmically combing through all the Go-related queries that cohort posted to the Bing search engine, the researchers

In the summer of 2016, the Pokémon Go app became a huge fad that got people all over the world walking around to capture Pokémon characters at virtual locations in their neighborhoods.

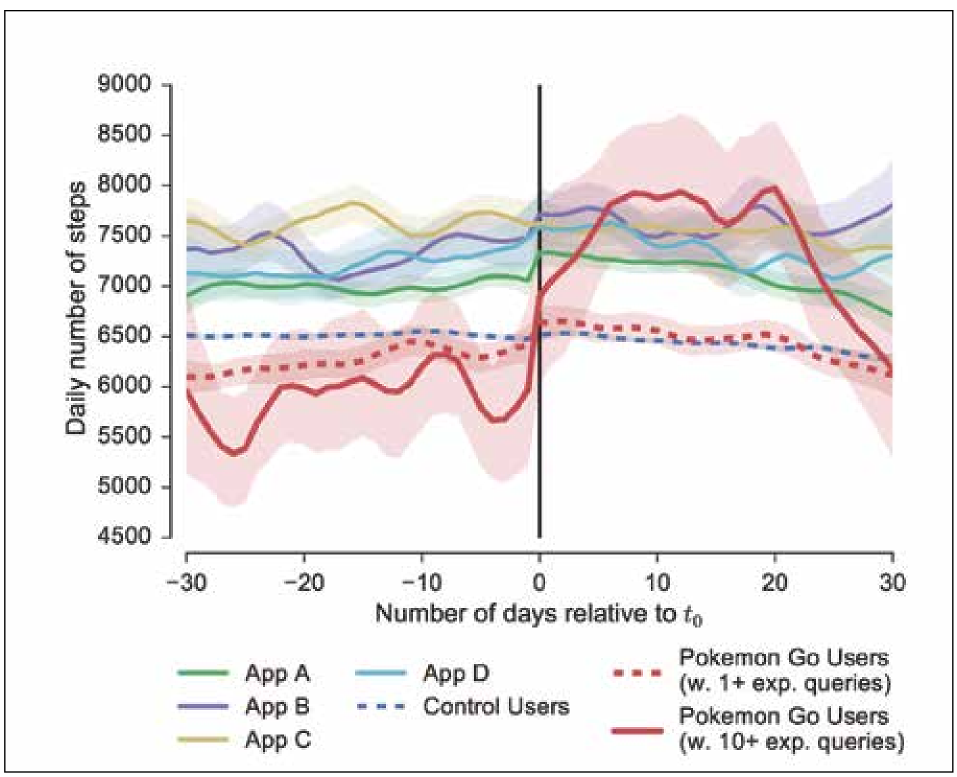

Mobilize Center researcher Tim Althoff and his colleagues at Microsoft showed that, before starting to play, Pokémon Go users are less active than average users of the leading consumer health apps (A, B, C, and D), but they experience larger increases in physical activity after starting to play (at t_0). To determine when a person was actually playing Pokémon, the researchers distinguished between search queries that suggested a user was merely seeking general information about the game, and "experiential queries" that indicated he or she was actively using it. Reprinted from Althoff T, White RW, Horvitz E, Influence of Pokémon Go on Physical Activity: Study and Implications, J Med Internet Res *2016;18(12):e315.*

inferred that just over 1,400 of those Band users were actively engaged with the game. And by examining users' Band data, they determined that Pokémon Go players walked as much as 1,473 extra steps a day—an increase of more than 26 percent over their prior activity levels.

Most of the people who use mobile health apps are already physically active. But Althoff's study suggested that even sedentary, obese, and older users benefited from playing Pokémon Go. And while the benefits of gameplay did fall off after three or four weeks, given the well-established link between physical activity and mortality risk, Althoff and his colleagues suggested that active engagement with the game—which has 65 million monthly active users—could nonetheless have a measurable impact on life expectancy.

The work also has immediate value to the research community. "The Pokémon Go project demonstrated how to do a physical activity study across more than 30,000 people with a very specific treatment," Althoff says—the treatment being playing Pokémon Go. "It shows how you can contextualize wearable data in a way that allows you to test large-scale interventions like this and provides a model for conducting such studies in the future."

In a pair of studies published this year, Althoff also exploited wearables to explore the real-world impact of participation in online social networks and app-based competitions.

Althoff and his co-authors, including **Jure Leskovec**, **PhD**, who spearheads the social and behavioral modeling effort at the Mobilize Center, analyzed data from Argus, a fitness-tracking app developed by the Silicon Valley startup Azumio. The app allows users to create posts about their physical activity (walking, cycling, yoga, etc.), and uses the accelerometers in their smartphones to unobtrusively track their physical activity.

In the first paper, Althoff wanted to see if participation in an online social network organized around fitness would affect physical activity in the real world. Argus provided the ideal data set: Althoff and his colleagues had access to anonymized information provided by 6 million users from 2011 to 2016—amounting to 10,000 times more users and a million times more activity tracking than most comparable studies—but the app's embedded social network was only added in 2013. This allowed the researchers to observe changes in physical activity among users who joined the network, and to compare their results to those of users who did not join.

Sure enough, people who made new social connections through the app increased their physical activity by approximately 7 percent, or 400 steps per day. Algorithms designed to tease out different kinds of effects, such as changes in internal motivation versus the influence exerted by new social connections, showed that 55 percent of the observed changes in user behavior were due to social influence. And Althoff and his colleagues developed a model that could predict which users would be most influenced by new social network connections—something that could contribute to the design of more effective apps in future.

In the second study, Althoff and his collaborators examined the impact of app-based fitness competitions on users' activity levels. They analyzed the data generated by 3,637 users who participated in 2,432 physical activity competitions over a 10-month period—again, the largest data set of its kind to date. And by considering factors such as age, gender, and prior activity level, the researchers were able to study which features of competition design were most likely to boost participants' activity levels. For example, competitions were most effective when participants shared similar levels of prior activity, and when there was a balanced mix of men and women.

According to Althoff, those kinds of insights could guide the creation of better app-based competitions. And that, in turn, would further his overall goal of optimizing online communities and mobile health apps "to help people be healthier and happier."

But the immediate impact of the project will be most acutely felt in the mHealth research community.

"Identifying social influence in observational network data is extremely challenging but very important for interventions," Althoff says. The Azumio project tested a new causal inference technique based on "delayed friendship acceptance," he says. "This worked really well and had never been done before."

A Comprehensive mHealth Platform and Specialized Sensors

Dunn and Althoff rely on commercially available apps and hardware to supply the data that fuels their algorithms. But the researchers behind MD2K—a Big Data to Knowledge Center of Excellence that brings together experts from 12 universities and the nonprofit startup Open mHealth—are taking a different route.

While they make use of some off-the-shelf wearables, they have also built their own novel sensors. In addition, they have developed an open-source mobile phone–based platform called mCerebrum, which includes more than 20 apps that combine and process the data from those various sources to support the discovery and validation of digital mobile-health biomarkers; and a cloud-based big-data component called Cerebral Cortex, which supports population-scale data analysis, visualization, and modeling.

According to **Timothy Hnat**, **PhD**, chief software architect for MD2K, mCerebrum can collect up to 70 million data samples per person per day, processing them in real time and running them through predictive models to trigger just-in-time health interventions. And because it's an open-source platform, it can be used by other researchers—with the potential to have a significant impact on mHealth research well beyond the BD2K program.

The mCerebrum sensor data is also synced to Cerebral Cortex, where

MD2K has developed mCerebrum, an open source, real-time software platform for data collection from sensors in smartphones and wearables. It can capture data from multiple sensors simultaneously while continuously evaluating data quality. It also allows for real-time data viewing, such as a live ECG signal. And it uses advanced analytics to convert the data into markers of health, behavior, and risk factors. The gyroscope data from a wrist sensor might, for example, provide insight into smoking or eating behaviors by revealing the telltale hand gestures involved in bringing a cigarette or piece of food to one's mouth. Images courtesy of MD2K.

health science researchers can visualize and interpret the information being gathered from study participants—and where data science researchers can perform large-scale machine-learning exercises to refine their algorithms.

Those refinements, in turn, trickle back down to mCerebrum, where the models that decide whether an intervention is required (and what kind of intervention it should be) can be fine-tuned on an individual basis, opening the door to data-driven personalized medicine.

Since the program's inception in 2015, researchers have used MD2K's best-of-breed wearables and software to track people's stress patterns and smoking behaviors, delivering alerts and behavioral exercises to help calm them or prevent them from taking a puff.

Now they are hoping to use a homegrown device called EasySense to more effectively treat congestive heart failure. This potentially fatal illness characterized by fluid buildup in the lungs affects almost 6 million Americans.

Unfortunately, says **William Abraham**, **MD**, the cardiologist at Ohio State University College of Medicine who leads the project, the standard method of managing congestive heart failure doesn't work very well.

Patients monitor their own symptoms and body weight (an indicator of fluid buildup), but rarely receive treatment until a significant amount of fluid has accumulated in their lungs—at which point "the horse is already out of the barn." Relapses are common, and readmission rates are higher than for any other cause of medical hospitalization.

EasySense could help change that.

The device, which is roughly the size of a hockey puck, emits pulses of ultra-wideband radio waves and listens to the echoes they create as they bounce off various bodily tissues. According to **Emre Ertin**, **PhD**, an electrical engineer at Ohio State who leads the development of MD2K's custom-built sensors, this allows EasySense to provide nearly EKG-quality heartbeat detection—and to gauge the fluid content of the lungs.

Abraham recently concluded a 20-person pilot study that successfully demonstrated the device could gather useful data in a hospital setting. He and his collaborators are now beginning a 75-person study in which participants will take the sensors home with them. By analyzing the data provided by EasySense along with the output from other wearables that record parameters such as respiration and oxygen saturation, Abraham hopes to determine which signals are most predictive of relapse and rehospitalization. A third and final study will then use that information to make treatment changes "to see whether or not we can actually keep patients out of the hospital," he says.

The goal, Abraham says, is to have mCerebrum send alerts and notifications directly to patients and their doctors before things get out of hand. They might, for example, suggest the need to reduce salt intake or prescribe an extra dose of diuretics when fluid levels begin to rise. As the data set grows, he expects patterns will emerge that will allow the team to tailor interventions on an individual basis.

Like the projects at the Mobilize Center, the work being done by Abraham and his MD2K colleagues promises to harness the Big Data generated by wearable biosensors to drive improvements in both personal and public health. And given the ever-increasing ubiquity of wearables (and the ever-increasing sophistication of data science), it's likely that the mobile health revolution is just getting started, spurred on by the methods and devices generated by the BD2K Centers—and available to the entire research community. □