

Putting Social Networks on Trial

There's a broad movement to use social media tools to find patients for clinical trials. But a host of technological and privacy hurdles present a challenge. *By Tracy Granzlyk Wetzel*

Recruitment of patients to fill clinical trials is a time consuming and project-limiting task required for the completion of a medical study. With advances in information technology being made across all areas of health care—from business practices in the C-suite to patient safety at the bed side, it would be logical to apply similar advances to the clinical trial process.

Patient social health networks


have grown in number over the last few years, and as a result, the amount of data collected about disease has been augmented in tandem.

For example, since its inception in 2004, PatientsLikeMe, based in Cambridge, Mass., has been trying to speed the process of bringing new therapies to market by connecting its patient community of more than 125,000 and giving context to the millions of data points these patients provide about the more than 1,000 conditions they

actively discuss.

PatientsLikeMe on one level is a social health network where patients can meet and share symptoms, treatments and experiences with others “like me”—struggling through a breast cancer diagnoses or living with multiple sclerosis, for example—says Chairman and Co-founder Jamie Heywood.

On another level, it's part of a health analytics infrastructure that can augment more tradi-



tional data collection methods.

“We are a health-outcome focused clinical discovery platform,” Heywood says. “Effective disease discovery requires patient quantification and stratification, and right now, medicine as a business practice does not do that. Electronic medical records do not really contain outcomes and we provide a way for patients to contribute their outcomes. Our business is ultimately in understanding the meaning of health and contributing to discovery.”

Heywood’s personal journey started in 1998 when his brother Stephen was diagnosed with amyotrophic lateral sclerosis (ALS). In a quest to find and develop a new therapy to save Stephen, Heywood, an engineer with a degree from the Massachusetts Institute of Technology, created the ALS Therapy Development Institute, a not-for-profit biotechnology firm focused on developing new ALS therapies. While successful in bringing two drugs to clinical trials, Heywood learned the complexities and costs inherent to the process and—all too painfully close to home—the time constraints related to success.

PatientsLikeMe, co-founded with his other brother, Ben, and friend Jeff Cole, also MIT engineers, is a result of those experiences. The company is one of a number of organizations try to use social networks and sophisticated screening and analytics technologies to create patient registries and feed clinical trials.

Leveraging communities

Clem McDonald, Ph.D., director of the National Library of Medicine’s Lister Hill Center for Biomedical Communication, oversees the *ClinicalTrials.gov* database and also sees advantages to the communities that social health networks provide.

“Recruitment of patients to clinical trials is important to the National Institutes of Health,” says McDonald. “Anything that could facilitate this would be very, very useful. It would speed the complicated process.” McDonald acknowledges that recruiting even 1,000 patients is slow going and typi-

cally takes much longer than expected.

“Currently, investigators can advertise in newspapers or on the Web, but in-ways are limited because of the fears around coercion,” he says.

The number of legal and privacy hoops that must be jumped through has yielded sophisticated matching systems to find the right fit for patients who are willing to participate in trials. For example, PatientsLikeMe is developing a system to enable patients in its database to be matched against the *ClinicalTrials.gov* database, a registry of both federally and privately supported clinical trials around the world.

According to Heywood, matches occur in real-time, as engineers download the entire *ClinicalTrials.gov* database each night. Approximately 75 percent of the PatientsLikeMe population can be matched to a trial; however, geography is a limitation, as only 50 percent of those who match to a trial can match to one within 75 miles of their home.

“Nearly half of our active members in the main disease communities on our site have accessed the trials list, and 60 percent of those have linked through to individual trials,” Heywood says. “We do not have data on the patients who actually enroll, but we when we push e-mails about trials, we accelerate the final study.”

From a technical point-of-view, matching patients to trials is a difficult process. Trials are often too specific or too broad in their definition of disease, making a match to exact criteria challenging.

“Because the trials are not coded by us,” says Heywood, “we have to match both specific and broader medical terms, which is a capability we are currently working on. The coding system at *ClinicalTrials.gov* is very flexible. It’s difficult to define eligibility because you need to do so by looking at who the patients are, and what the data architecture is that will make it all work.”

The National Library of Medicine (NLM) provides the keys for mapping trials into their system by computing the medical subject heading (MeSH) codes

from the investigator-entered trial descriptions and includes the codes in the machine-readable clinical trial data files. PatientsLikeMe then uses NLM’s Unified Medical Language System (UMLS) to map the MeSH codes onto the SNOMED hierarchy, which the health data integrity team has already coded with PatientsLikeMe conditions. They are still searching for “the perfect” MeSH-to-SNOMED mapping, as matches are missed when coding is either more or less specific in SNOMED. The future calls for greater flexibility in searching the ontology, which will improve the precision of the matching process.

A need for speed

Speed is at a premium during the matching process because time is an enemy for getting a therapy to market. The Pharmaceutical Research and Manufacturers Association (PhRMA) states that the drug development process takes anywhere from 10 to 15 years, from pre-discovery through the FDA review process. The completion of related clinical trials can eat up six to seven years of that time frame. And with an average cost for the journey estimated to be between \$800 million to \$1 billion to manufacturers, which is undoubtedly passed on through the health care system, it would seem that inroads to ethically shorten the process through technological advances would only make sense.

Paul Harris, Ph.D., director at the Office of Research Informatics at Vanderbilt University, is developing new ways to match patients to clinical studies. The organizations disease-neutral, Web-based recruitment registry, ResearchMatch, comprises more than 19,000 individuals interested in learning more about clinical trials, along with 991 researchers at 65 institutions recruiting for 302 active studies.

According to Harris and colleagues, the challenges relative to patient recruitment for both clinical and translational research impedes the entire process. And with approximately 75 percent of the gen-

“The idea is to get science from bench to bedside.”

—Paul Harris

eral public reporting a lack of knowledge about clinical trials in general, and only 2 percent of the U.S. population participating in a clinical trial on an annual basis, it would stand to reason that the ability to cast a wider net of information and opportunity is an important endeavor.

“The idea is to get science from bench to bedside as quickly as possible,” says Harris. “In my role, I enjoy the opportunity to look across the entire research enterprise and assemble novel ways to use informatics to increase the throughput of translational research.”

While the ResearchMatch registry was conceptualized and developed by Harris’ team at Vanderbilt, it is now a collaborative effort with the Clinical and Translational Science Awards institutions, of which there are 60 across the country, and funded by the National Institutes of Health.

Anyone in the U.S. 18 years of age and older can go to the site and volunteer to participate in ResearchMatch. Researchers from affiliated medical centers or hospitals in the network, along with their institutional liaisons, are the second primary stakeholder included in the registry. The software and back-end infrastructure are built on a PHP and MySQL framework. PHP is a scripting language used or developing Web applications, and MySQL is an open-source database built on a relational database management system.

Volunteers create a unique user name and password, and then provide demographic, geographic and relevant health information about conditions and medications, which is structured to then prompt the UMLS metathesaurus keyword-matching algorithm built into the system.

“Once registered,” says Harris, “volunteers may be considered and contacted for many types of trials across the United States. Potential volunteers may also click-through to public information on diseases provided by sites such as MedlinePlus, *ClinicalTrials.gov* and NetWellness. Researchers at participating institutions can search for potential study patients based on their exclusion and inclusion criteria, and then send an IRB approved message to those appearing to be good candidates.” All volunteer data is deidentified until that volunteer gives the OK by clicking on hypertext links confirming their interest, and requesting more information from the research center.

Concerns around a project like ResearchMatch include the need for privacy and data protection, as well as approval from the Institutional Review Board, the organization that also oversees and monitors U.S. trials.

“There are legal agreements that need to be signed to join the network of participating institutions, and then those institutions have to assign someone onsite to police it, and work with their local IRB so that everyone is comfortable with the recruiting process,” Harris says.

Looking local

In addition to ResearchMatch, Harris’ team also has capabilities to recruit local Vanderbilt employees and patients who enter the Vanderbilt hospital system. By linking to their human resources database and EHR within the health system, they have been able to further expand their reach, as well as decrease the workload of research coordinators.

“From an informatics perspective, we are helping find patients in various ways,” shares Harris. “We have created an internal listserv, and with IRB approval, can send informative e-mails [about studies] to all employees. If they choose to opt out, they will not receive e-mails like this again. Our research team tells us this has been very effective.”

By monitoring the EHR, Harris’ team can detect signals that may identify a current patient as a match for an open trial. The research coordinator and the attending physician caring for the patient can be notified, as well as the patient in certain circumstances.

According to Harris, this all revolves around the way rules are written to define exclusion/inclusion criteria and derive triggers when the signal is strong enough, and the criteria have been met to initiate a notification mechanism.

“We have been experimenting in notifying patients directly about individual studies through a patient health portal,” says Harris, “but here we have to be sensitive to patient perceptions and sensitivities concerning medical research.”

According to Harris, writing a small class of rules based off information found in the EHR, such as previous clinics visited, ICD-9 codes, medications, or common demographic elements, computer algorithms can greatly reduce the “haystack” of individuals who might otherwise be considered for potential contact about research studies.

“In some cases, we’re reducing 95 percent to 99 percent of the workload that a research coordinator might do by asking the computer to pre-filter and return a much smaller subset of potential patients for them to sort through,” shares Harris. “In this case, we’re focusing on developing generalized methods to assist many research studies, rather than focusing on specificity for a handful of studies.”

Mark Pletcher, M.D., associate professor of epidemiology and biostatistics at the University of California-San Francisco, is contributing to the cause via a pilot study. As principal investigator, he is examining the use of social networks to recruit large numbers of people in order to facilitate both observational and interventional studies on the prevention of cardiovascular disease.

“Our primary goal is to facilitate and lower the cost of doing interventional research which is incredibly expensive but necessary,” says Pletcher. “If we can figure

out ways to reduce the marginal cost of conducting studies this would be a huge contribution.”

Web recruiting

Pletcher and his team started recruiting on Dec. 1, 2011 at www.socialheartstudy.org, using the Web site to “pilot” the study and see how far they can get with recruitment through their own social networks, without spending money to advertise. As of Jan. 9, 306 people had gone through the baseline questionnaire. The goal for the pilot is 1,000 participants and once complete, they will seek funding for a larger scale study. The long-term target is 1 million unique participants. The only inclusion criteria are a valid e-mail address and being 18 years of age or older.

“One of our goals is to prevent cardiovascular disease by giving participants a personalized cardiovascular health profile based on information they provide. Each participants’ personal profile shows their overall short-term and lifetime risk of cardiovascular disease, and flags which, if any, of their cardiovascular risk factors are non-optimal,” Pletcher says. “As the study develops, we’d like to continue this strategy. For every additional piece of data we collect from participants, we’ll try to provide some analysis. For example, we plan to offer a detailed dietary questionnaire, and then offer a personalized dietary profile describing how the participant’s diet differs from a healthy diet, and from the diets of others within and outside of their social network”

The secondary goal of the Social Heart Study is to examine how social networks impact an individual’s overall health. Co-Principal Investigator, James Fowler, Ph.D., professor of medical genetics and political science at the University of California-San Diego, has shown health behavioral characteristics tend to spread through social networks, and will now explore if the same is true of online social networks. This study encourages participants to link to their Facebook page through Facebook Connect, allowing researchers to then ask ques-

tions such as: Do Facebook users with more friends have better cardiovascular health? Are people who spend lots of time on Facebook healthier than people who spend less time? Is cardiovascular health transmitted through friend networks, and if so, if your friend becomes healthier, do you?

“We have been dreaming about this study for a couple years now,” says Pletcher. “The intent is to capture a core set of cardiovascular health measurements including blood pressure, diabetes, smoking, and cholesterol levels in this cohort over time, add on additional measures such as diet, exercise, sleep habits for interested participants, and then conduct interventional studies to try to improve health.”

Pletcher’s vision for the future includes adding special measurements captured by external sensors such as step counters or even iPhone apps to facilitate interventional studies and education that could be delivered through social media.

Online social health networks are growing both organically and with intent across the Web, many with no other purpose than to provide much needed social support to patients. However, organizations such as PatientsLikeMe and Alliance Health Networks, based in Salt Lake City, are trying to harness the power these networks hold, both as business opportunities and the impact they can have on changing health care.

Diabetic Connect (www.diabeticconnect.com) is the largest of 52 disease-based communities under the management of Alliance Health Networks, established in 2006. In December, Alliance reported just over 1 million patient profiles, with 500,000 of those from the Diabetic Connect community. The company’s mission is to create the best user experience and facilitate the social and sharing aspects of their communities.

“It’s often helpful to hear about others’ experience with their doctor, and how a good relationship with a doctor works,” Crowley says. “There is a shortage of diabetes specialists, for example, and so

many are diagnosed by a general practitioner that doesn’t specialize in managing the disease. A network like this helps broaden a patient’s understanding.”

Alliance is currently in discussions and working with partners such as Clinical Research Organizations (CROs) and patient recruitment firms who are interested in helping connect patients in their communities to clinical trials.

Treading cautiously

“We have done some things to help patients find trials—recruiters have approached us and we have created banners on the site for members to learn more as a way to pre-qualify people,” he says. “We are very careful about privacy. The pilot programs we’ve done have been “opt in” where patients would have to request more information.”

All data is encrypted and securely stored, and patient profiles cannot be crawled by search engines. Crowley says the company does not aggregate data or provide patient data to study sites in any form unless the patients themselves have requested it.

Joey D’Alessandro, vice president of client services at Alliance, works on the business-to-business side, and shares that they have contracted with CROs and patient recruitment firms to identify qualified referrals by driving traffic to a client’s screening form.

Having worked at a CRO prior to Alliance, D’Alessandro is aware of the challenges inherent to clinical trials—many trials are narrow in their design and limit those who can participate by geography, demographics and medical conditions.

“Studies are often times designed in silos and target the ‘unicorn,’” he says. “The study design may be great, but you will never find enough of those individuals to make it statistically significant yet they expect someone to go out and find those unicorns.” ■

Tracy Granzyk Wetzel is a Chicago-based health care freelance writer.