



BRATISLAV MILENKOVIC

Hours before the close of Kaggle's competition to find out why almost one-third of women in the United States are not screened for cervical cancer, the leading team has submitted the 115th iteration of its model. Forty groups around the world are competing to win US\$100,000 in a challenge sponsored by biotechnology company Genentech.

The models are based on analyses of a 150 gigabyte database of de-identified patient data, says computational biologist Wendy Kan, who set up the challenge and works at Kaggle in San Francisco, California, a company that runs predictive modelling and analytics competitions that allow data scientists to compete to solve complex problems. In addition to finding solutions, contestants are asked to explain their reasoning. "It's very important for us to tell a story," Kan says. Later, on a Kaggle forum, a member of the winning team presents two of the group's hypotheses: multiple chronic diseases and mental-health issues are major factors in why some women skip screening.

Another Kaggle challenge, which began in December, asked participants to transform the diagnosis of heart disease by coming up with an algorithm to examine cardiac magnetic resonance imaging (MRI) scans to see how well the heart is pumping blood — "A very difficult problem," Kan says. Entrants used a cardiac MRI data set provided by the US National Heart, Lung and Blood Institute, and 192 teams were in the running for the \$200,000 prize when the competition closed. The victors were two quantitative analysts who have worked with hedge funds, but had no experience in cardiology.

So far, more than 450,000 data scientists have tried their hand at Kaggle's predictive-modelling puzzles, says economist Anthony Goldbloom, founder and chief executive of the organization. The problems — many pertaining to health, but others in fields that range from criminology to search technology — are set up so that the background of entrants doesn't matter, he says. As long as they have suitable modelling skills, no particular experience or qualifications are needed.

"They are all smart, highly motivated and incredibly capable," adds Goldbloom. "The winning margin is usually very small; often the difference between first and second isn't even statistically significant."

Kaggle is one of a number of organizations running open global challenges in life sciences to address knotty problems in basic biology, clinical research or health care. The approach is steadily gaining backers in academic laboratories and classrooms, drug companies and government agencies as a way to bring well-defined, but thorny problems to the attention of brilliant minds around the world.

The design of the competitions varies from challenge to challenge and host to host. Some ask for modelling algorithms, others for ideas, and still more for prototype medical solutions.

CHALLENGES

# Crowdsourced solutions

*Open competitions bring new minds, skills and collaborations to problems in biomedical research.*

BY ERIC BENDER

Prizes are often offered, although participants usually insist that money is not the main motivation. Some of the winning solutions, especially those sponsored by industry, remain secret, but others are made openly available and a few have already resulted in advances in clinical research.

### CLOCKWORK ORIGINS

Competitions in science and engineering have a long history. In 1714, the Longitude Act saw the UK government offer a reward of £20,000 (well over £2,000,000 (\$2,865,400) in today's money) for a solution to the problem of calculating longitude at sea. Not just one, but two answers emerged: the marine chronometer, developed by clockmaker John Harrison, which kept time at sea well enough for navigators to calculate longitude effectively; and a method for devising longitude from the motion of the Moon borne of a combined effort by scientists, including mathematician John Hadley and astronomer Tobias Mayer.

But it took the advent of the Internet for crowdsourced medical contests to really take off, notably with the Critical Assessment of Protein Structure Prediction (CASP) experiments, which have seen research groups test their methods for predicting 3D protein structures against those of their peers since 1994.

The competitions have gained more industry backing as pharmaceutical companies began to struggle with their pipelines. The crowdsourcing firm InnoCentive, for example, formed in 2001, a time when “the pharmaceutical industry needed to rethink its business model”, recalls Alph Bingham, co-founder of the company based in Waltham, Massachusetts, and then a vice president at pharmaceutical giant Eli Lilly. “The Internet let you access minds on a scale and a scope that had never been possible before.”

Spun out from Eli Lilly, InnoCentive has held more than 2,000 open challenges and attracted more than 375,000 ‘solvers’. The continual string of challenges can be tightly focused and relatively small, such as a \$30,000 challenge to find a minimally invasive skin-biopsy method to measure gene expression, or attempt to tackle larger problems, such as a major \$500,000 challenge sponsored by the US National Institutes of Health (NIH) to look for robust methods to examine individual cells. Proposals such as these are inherently risky and might not survive the conventional NIH grant process.

Indeed, challenges seem to hold a number of advantages over conventional research practices. One of the leading crowdsourcing initiatives is the Dialogue for Reverse Engineering Assessments and Methods (DREAM) Challenges programme, which sees groups compete in open competitions to solve complex

**“Prizes can really bring a new population of researchers into the field.”**



Scientists from around the world competed to win a BioMed X fellowship in Heidelberg, Germany.

modelling problems in systems biology, says Gustavo Stolovitzky, co-founder of the project and computational biologist at IBM's Thomas J. Watson Research Center in Yorktown Heights, New York.

When dozens of teams around the world take on a DREAM project, they often accomplish in months what would take a single research group years, “since you can multiply the number of people working on the problem by 50 or 100,” says Stolovitzky. Many challenges also bring in researchers from other fields, who may approach problems in ways that those closely acquainted with them would not.

Just as crucially, challenges jump-start collaborative communities. For instance, the ICGC-TCGA DREAM Somatic Mutation Calling Meta-pipeline Challenge is a collaboration between DREAM, the International Cancer Genome Consortium, The Cancer Genome Atlas and biomedical research organization Sage Bionetworks in Seattle, Washington. Its aim is to improve standard methods for identifying cancer-associated mutations and rearrangements in whole-genome sequencing. In the process, they are building an ongoing community in which researchers can find the best and latest algorithms, rather than having to go to scientific journals.

Crowdsourced tournaments can also open up access to data — either those aggregated specifically for the purpose, such as Kaggle's cervical-cancer and cardiac MRI databases, or data sets that would otherwise lie dormant. “There are too many data silos in which researchers hoard their data, sometimes for years,” Stolovitzky says. “Ultimately, everybody should be able to look at that data with information about how the data was gathered, allowing collaboration and data

sharing in a positive and meaningful way.”

In addition, contests can lower the legal barriers that plague collaborations between institutions or companies, says Bingham. “They offer ways to engage all these different people without having to precede that whole process with 200 days of legal briefs being exchanged between institutions,” he says.

For these contests to achieve these positive impacts, however, they have to be well managed. Crowdsourcing is of little help in areas in which research is at such an early stage that the organizers can't ask the right questions. For any challenge to work, the problem must be well-defined and able to be judged fairly, says systems biologist Stephen Friend, co-founder and director of Sage Bionetworks. It's also important for an impartial expert in the field to act as a convener and nurture the emerging community, he says.

Non-profit foundations — increasingly important providers of research funding — are also making use of crowdsourcing. Often these focus on diseases that drug companies rarely target (see page S68). One example is Prize4Life in Berkeley, California, founded in 2006 when Harvard business school graduate Avichai Kremer was diagnosed with amyotrophic lateral sclerosis (ALS; also known as motor neuron disease), and best known for its \$1-million contests.

“Prizes can really bring a new population of researchers into the field,” says neuroscientist Neta Zach, chief scientific officer at Prize4Life. “And a lot of them continue to work on ALS.” Prize4Life's first major challenge addressed the lack of useful biomarkers for ALS progression. “We expected that the tool would be based on measurements from blood or cerebral spinal fluid,” Zach says.



Participants at a Massachusetts Institute of Technology Grand Hack discuss health-care challenges.

JEANETTE CAJIDE

Instead, the winning tool in 2011 was a more creative solution: a pain-free non-invasive medical device that measures the flow of electrical current through muscle tissue. The winnings helped to build the San Francisco start-up Skulpt, which is testing such devices in ALS trials (as well as offering them to consumers as fitness tools).

The foundation also partnered with DREAM and InnoCentive in a \$50,000 challenge to predict the progression of ALS. When the predictions of the winning algorithm were compared with those made by ALS clinicians in the assessment of 14 people with ALS (R. Kuffner *et al. Nature Biotechnol.* 33, 51–57; 2015), “the algorithm outperformed each and every one of the clinicians on each and every one of the patients”, Zach says. The model is now used to make ALS clinical trials more efficient and their results clearer — a better understanding of ALS makes it easier to assess the benefits of treatment.

DREAM was launched in 2006 by Stolovitzky and systems biologist Andrea Califano at Columbia University in New York City to improve the state of the art in systems-biology modelling. As well as solving problems, DREAM challenges validate the solutions.

Sometimes when data-science groups tackle a difficult problem, they can convince themselves that they have produced a good solution, rather than actually solving it well. Stolovitzky calls this the “self-assessment trap”, which can lead to mistakes such as overfitting models to one set of data. But if 50 DREAM teams are involved, “we can see if we can really find a clear signal in the data”, he says.

In 2012, DREAM joined forces with Sage Bionetworks, which had created Synapse, a pioneering open-computing platform for data analysis and sharing. The first joint challenge generated models to classify the aggressiveness of breast cancer. The models clearly performed better than today’s commercial tests, says Friend. “More importantly, the challenge showed that people who had not generated the

data were able to get deep insights,” he says. “And the electrical engineer who won had very little chemical background.”

### RISE TO THE CHALLENGE

Competitions are beginning to exploit the opportunities provided by data contributed directly by patients. Sage, for example, created mPower, an app that uses iPhone sensors to measure symptoms of Parkinson’s disease progression such as dexterity or gait. And Sage has partnered with other groups, such as Oregon Health and Science University in Portland and Harvard University in Cambridge, Massachusetts, to create numerous such apps, which

**“At the end of the day, cash is often a scorecard, not a paycheck.”**

can very quickly provide high-quality data. “We have over 200,000 people who have said, I want to share my data with qualified users,” Friend says.

In November 2015, a DREAM hackathon drew participants for two evenings of pizza, beer and the opportunity to begin interpreting data from tens of thousands of mPower users. That event reflects another trend in crowdsourcing — the rapid spread of biomedical hackathons. These are designed to bring experts from different disciplines face to face. The Hacking Medicine initiative at the Massachusetts Institute of Technology (MIT) in Cambridge, for instance, has so far hosted almost 50 such events, teaming up engineers and data scientists with clinicians in 1- or 2-day events that are meant to quickly and iteratively work towards initial solutions to a host of health-care problems.

Among early results is an infant-resuscitation device for use in developing countries. The Ugandan paediatrician who first presented the problem has now taken the device into clinical trials in his country. The MIT initiative has helped to spark similar gatherings in places such as India and Uganda, led by the Consortium for Affordable Medical Technologies at

Massachusetts General Hospital in Boston.

Bringing researchers with varied expertise and skills together in one physical location can accelerate research. The BioMed X Innovation Center in Heidelberg, Germany, has gone further with what co-director, and biologist, Christian Tidona describes as an “outcubator”. Researchers compete not to come up with the best solution, but for the chance to try.

BioMed X begins by posting a very specific problem from one of its sponsors online. This could be exploring a new drug target or an area of treatment new to the sponsor. These requests typically get 400–600 responses from around the world. BioMed X picks 15 of the most promising concepts submitted and brings their creators to Heidelberg, where they form teams for an intense 5-day competition. The winning group then tackles the problem in two- to four-year fellowships in Heidelberg.

One of the first teams to go through the four-year exercise — made up of researchers from Germany, Slovenia and Egypt — created bioinformatics tools for designing highly selective inhibitors of kinases, proteins that play a part in many diseases. The sponsor, Merck, bought the intellectual-property rights and then licensed them back to the team, which formed a start-up company to develop the technology.

### RULES FOR THE FIGHT

The benefits for research are clear, but what is it that drives participation in crowdsourced competitions? When a challenge is centred in a researcher’s field, typically the greatest incentives to participate are the chance to publish a paper in a top journal and to network with peers, organizers say.

But often the entrants are not the usual suspects. “They’re also gadgeteers, basement inventors and weekend engineers,” says Bingham. “It’s not a bunch of French-literature majors that are solving our chemistry problems, but it might be physicists or intellectual-property attorneys or biologists.” Even in competitions with cash prizes, “at the end of the day, cash is often a scorecard, not a paycheck”, he says. Challenges would be “a silly way to make money”, says Goldbloom. The main draw for participants is what originally led him to found Kaggle — the desire for “access to interesting data sets and interesting problems”.

For medical firms, the challenges often provide a relatively quick and inexpensive way to solve tricky problems, Bingham says. At the same time, he points out, “in order to bring a product to market, they usually have to solve a thousand problems of equal complexity”. For all concerned, “the wisdom of crowds works beautifully in a great percentage of the cases”, says Stolovitzky. “We’re seeing a lot more buy-in for these challenges. If you can multiply the number of people, you can accelerate the research.” ■

**Eric Bender** is a science writer based in Newton, Massachusetts.