CROWDSOURCING TO EXPERTS:
The DREAM Challenges

By Esther Landhuis

At a 1906 county fair in England, some 800 villagers tried to estimate the weight of an ox. None of the contestants hit the mark, but a closer look at their guess cards led to a stunning discovery. Stacking the estimates from lowest to highest, the middlemost value came within 0.8 percent of the ox's butchered weight—closer than individual guesses submitted by cattle experts. Published in 1907, these findings on the statistical concept of median were among the first to demonstrate the wisdom of crowds.

more than 30 open science competitions drawing diverse experts to complex biomedical questions. Wooed by prize money and opportunities to publish their approaches in top journals, researchers around the globe have developed computational models for a variety of translational medicine challenges, including predicting drug responses and disease outcomes.

DREAMing of Better Solutions

For systems biology, the crowdsourcing concept emerged as scientists were faced

© Dwnld777 | Dreamstime.com

A century later, leaders of the genomic revolution are summoning "crowds" to tackle some of the toughest problems in modern medicine. These aren't crowds of ordinary townsfolk—or even biologists, necessarily. Many train in fields such as computer science, engineering or statistics and spend far more time staring at numbers and graphs than scrutinizing cells under a microscope. They're part of a collaborative initiative called DREAM (Dialogue on Reverse Engineering Assessment and Methods). Since 2007, the group has organized

with organizing huge piles of data coming out of DNA microarray experiments. Microarrays measure the expression of thousands of genes at once, comparing their levels in groups of cells under normal versus disease conditions, for instance. But massive lists of differentially expressed genes by themselves aren't that useful. Researchers want to understand how the genes are connected, such as whether they encode proteins that interact or regulate other genes, says computational biologist **Gustavo Stolovitzky**, **PhD**, of IBM Research and the

> ...Leaders of the genomic revolution are summoning "crowds" to tackle some of the toughest problems in modern medicine. These aren't crowds of ordinary townsfolk— or even biologists, necessarily. Many train in fields such as computer science, engineering or statistics and spend far more time staring at numbers and graphs than scrutinizing cells under a microscope.

Icahn School of Medicine at Mount Sinai in New York, one of DREAM's founders.

Computational scientists have assembled networks using algorithms that reverse-engineer or infer gene relationships from data. However, some worry that validating these approaches relies too much on cherry-picking. By focusing on connections that seem consistent with prior publications, "you're selecting what works for you but forgetting the ones that might not be working," Stolovitzky says.

DREAM originated as a way to evaluate these network inference algorithms. Open competitions allow participants to see which schemes work and which don't. Soon, the group realized DREAM challenges could do more than assess methods—they could ac-

only their genomic, epigenomic and proteomic profiles.

The 44 algorithms submitted by the research community suggest it is possible to develop rational approaches for predicting drug responses. However, their predictions are "not yet as good as we would like," Stolovitzky says. Asked to rank cell lines from most to least sensitive for each drug tested, the top model ordered 60 percent of cell line pairs correctly. By comparison, "a monkey doing this task would be right half the time," notes Stolovitzky.

In a related DREAM challenge, participants devised algorithms to rank 91 pairs of compounds on how strongly they enhance or sabotage each other's effects—otherwise known as synergism and antagonism. This

treatment, and the disease course varies widely between individuals. Most patients die three to five years after symptoms appear, but some make it 10 years past onset.

Disease variability is a big challenge for the field, says neurologist **Merit Cudkowicz, MD, MSc**, an ALS specialist at Massachusetts General Hospital in Boston. It means clinical trials need to be large for tested compounds to show an effect. The heterogeneity also suggests different biological mechanisms could be at work in patients who decline more quickly or slowly. So "maybe there will be therapies that work in some people but not in others," Cudkowicz says.

Challenge organizers supplied competitors with three months of lab test data as well as demographics and family history for

By focusing a community of experts on a specific problem for a limited time, work that might take 10 years in a single lab could be done by the crowd in several months. Reliability also got a boost. "When we aggregate the solutions from all participants, the resulting solution is often better than the best," Stolovitzky says.

celerate research. By focusing a community of experts on a specific problem for a limited time, work that might take 10 years in a single lab could be done by the crowd in several months. Reliability also got a boost. "When we aggregate the solutions from all participants, the resulting solution is often better than the best," Stolovitzky says.

Challenge:
Predict Cancer Drug Responses

Several papers published last year in *Nature Biotechnology* highlight DREAM challenges aimed at developing rational approaches to predict how cancer patients respond to treatments. These days, choosing drugs involves a fair amount of guesswork, unless the patient happens to have a gene mutation known to drive that particular cancer. In one challenge, the DREAM coordinators gave teams genomic, epigenomic and proteomic profiles for 35 breast cancer cell lines as well as information on how the cells respond to treatment with a group of drugs. The teams were then asked to predict how well a different set of 18 cell lines would respond to those drugs, given

challenge proved harder—only 3 of 31 submissions performed better than random guesses. However, the top methods were based on different hypotheses about how synergism and antagonism work—and combining them produced better results and provided insights into how drug interactions might work.

Thus, "while the results are not immediately applicable to the clinic, they begin to establish the rules and types of data needed to predict accurately the correct drug regimen," says **Dan Gallahan**, **PhD**, who directs cancer biology research at the National Institutes of Health (NIH) in Bethesda, Maryland. "This is the type of research needed to make precision medicine a reality."

Challenge:
Predict Neurodegenerative Disease Progression

One of the more successful DREAM initiatives offered $50,000 for the computational approach that could most accurately predict disease progression in people with amyotrophic lateral sclerosis (ALS). This neurodegenerative disorder has no effective

1,822 people enrolled in ALS clinical trials. DREAM teams were then asked to predict each participant's disease progression over the subsequent nine months.

The ALS challenge drew 1,073 registrants from 64 countries. Top-performing models predicted disease outcomes better than a panel of 12 experts, and the winners "didn't know anything about ALS," Stolovitzky says.

Statisticians estimate that the best two algorithms could reduce the size of ALS clinical trials by 20 percent. For a 1,000-patient Phase 3 trial, that would save $6 million. One company—Origent Data Sciences in Vienna, Virginia—is working to incorporate new predictive analytics into future ALS trials. By estimating how a patient's symptoms would progress without the intervention, these tools are particularly useful in early trials that lack placebo arms, says Origent CEO **Mike Keymer**.

Researchers won't know the true impact of DREAM algorithms for a while. But in the meantime, the challenges have succeeded in getting cross-disciplinary researchers out of their silos and working together. □