

PHILOSOPHY OF SCIENCE

Machine Science

James Evans and Andrey Rzhetsky

Scientists today cannot hope to manually track all of the published science relevant to their work. A cancer biologist, for instance, can find more than 2 million relevant papers in the PubMed archive, more than 200 million Web pages with a Google search, and databases holding results from experiments that produce millions of gigabytes of data.


This explosion of knowledge is changing the landscape of science. Computers already play an important role in helping scientists store, manipulate, and analyze data. New capabilities, however, are extending the reach of computers from analysis to hypothesis. Drawing on approaches from artificial intelligence, computer programs increasingly are able to integrate published knowledge with experimental data, search for patterns and logical relations, and enable new hypotheses to emerge with little human intervention. Scientists have used such computational approaches to repurpose drugs, functionally characterize genes, identify elements of cellular biochemical pathways, and highlight essential breaches of logic and inconsistency in scientific understanding. We predict that within a decade, even more powerful tools will enable automated, high-volume hypothesis generation to guide high-throughput experiments in biomedicine, chemistry, physics, and even the social sciences (1).

Proponents of data-driven science (2–4) conjecture that hypotheses are obsolete: New knowledge will simply emerge from mechanical application of algorithms that mine data for plausible patterns. This approach is attractive, but there are potential pitfalls. The discovery of patterns from data alone is similar to the task faced by an explorer in an unfamiliar jungle, without a guide. With no sense of what is already known about the environment or its perils, she is likely to misclassify what she sees—fearing the intimidating but harmless snake; ignoring the tiny lethal frog.

Recent research demonstrates how scientists can use computers to become better-informed and more agile explorers. New computational tools can expand the pool of concepts and relations used for generating automated hypotheses by (i) drawing more

from the vast corpus of published science, and (ii) synthesizing new higher- and lower-order concepts and relations from the existing pool of knowledge. This approach can enable scientists studying a particular natural system, such as a biochemical pathway, to identify and fill in missing pieces, and traverse reasoning chains much longer than those possible with the unaided mind. For example, researchers have used computation to increase the number of candidate genetic aberrations considered in synthesizing hypotheses about disease (5–7). They have also increased the number of potential biological activities involved in describing new gene functions (8, 9) and ironed out past errors (10). Similarly, scientists have used computation to increase

Soon, computers could generate many useful hypotheses with little help from humans.



Logical leaps. Scientific knowledge and concepts can be represented as jigsaw puzzle pieces that, with the help of new computational tools, can be assembled into new hypotheses. In Swanson’s ABC model, if the literature from one scientific subfield includes two concepts (A, red, and B, yellow), and the literature from another subfield includes B and C (blue), then an analyst may computationally infer that A and C are directly or indirectly related, potentially leading to new hypotheses that cross subfield boundaries.

the potential number of proteins and metabolites involved in biochemical networks, and to generate predictions about which locations in those networks could be altered to improve health (11) and to identify elements misidentified as participating in a network (12).

Merely increasing the pool of concepts and relations, however, would simply generate multitudes of low-quality hypotheses. Scientists can profitably restrict that multitude by using a selection process that draws on insights into the social, cultural, and cognitive production of science. For example, Swanson pioneered the ABC model of hypothesis generation, which focuses on hypotheses that cross boundaries between distinct scientific literatures. If concepts A and B are studied in one literature, and B and C in another, Swanson assumed transitivity to hypothesize that A implies C (see the figure and fig. S1). He then demonstrated that novel A-to-C inferences were likely to be true, although unlikely to be arrived at via other means (13–16). Through this approach, Swanson hypothesized that fish oil could lessen the symptoms of Ray-

naud’s blood disorder and that magnesium deficits are linked to migraine headaches. This heuristic relies on an implicit understanding of scientific communities and publishing norms. It assumes that unpublished ideas within a research community are less valuable than ideas that link seemingly unrelated communities. Within a subfield, scientists are typically familiar with all of “their own” ideas, so unpublished connections more likely represent negative knowledge—superficially plausible ideas that participants know are wrong from experience. Unpublished ideas about subjects (such as the role of particular molecules or genes) that cross subfield boundaries, however, are much more likely to represent unasked questions. A recent analy-

sis of biomolecules common to several fields of biomedicine, for instance, suggests that many communities could profit from generating predictions that bridge field boundaries and link disparate properties of these molecules or other scientific concepts. (17).

Automated expansion of concepts and relations across community boundaries is severely constrained by incompatibilities in the language used by different scientific communities (18, 19). Because subfields have distinct histories, they often use different language to express the same concept, or similar terms to refer to unrelated entities. If researchers could computationally map these languages onto one another, as some are beginning to do with medical terminologies (20), they could vastly increase the number of possible hypotheses. Mapping concepts across languages would highlight parallels in theories from different domains, as well as changes in meaning with time (semantic drift) and multiple meanings (polysemy). These differences could be computationally mined to identify novel conceptual linkages.

By prioritizing hypotheses that contain concepts spanning existing scientific theories, languages, and cultures, investigators could productively focus on the most novel (21).

Analysts can also increase the pool of concepts and relations by computationally synthesizing new concepts and relations from those previously published. Computers have been deployed to “coarse grain” or identify new, higher-order aggregates of established concepts within studies of biological pathways, medical syndromes, and social classes (22). Scientists have also discovered new relations by identifying regular similarities between existing elements (23, 24). They have efficiently constricted the vast number of possible new aggregates by focusing on those that share physical properties or patterns, or integrate components of a broader system, such as a particular disease.

In the past, computational approaches have been more successful in small, well-

defined systems than in larger, less studied, or more complex ones. The explosion of data from high-throughput experiments, however, increasingly presents researchers with very complicated systems. Facing these data with questions equal in scale and complexity will be critical because, in the words of Mark Twain, “you can’t depend on your eyes when your imagination is out of focus” (25).

References and Notes

1. D. Lazer *et al.*, *Science* **323**, 721 (2009).
2. L. Hood, *BioCommerce Week*, 9 September 2004.
3. T. Golub, *Nature* **464**, 679 (2010).
4. M. A. O’Malley, K. C. Elliott, C. Haufe, R. M. Burian, *Cell* **138**, 611 (2009).
5. C. Perez-Iratxeta, P. Bork, M. A. Andrade, *Nat. Genet.* **31**, 316 (2002).
6. I. Iossifov, T. Zheng, M. Baron, T. C. Gilliam, A. Rzhetsky, *Genome Res.* **18**, 1150 (2008).
7. J. Liu *et al.*, *Science* **323**, 1218 (2009).
8. R. D. King *et al.*, *Nature* **427**, 247 (2004).
9. R. D. King *et al.*, *Science* **324**, 85 (2009).
10. J. Blake, personal interview, 5 February 2009.
11. D. A. Ruths *et al.*, *J. Comput. Biol.* **13**, 1546 (2006).
12. T. L. Hsiao, *et al.*, *Nat. Chem. Biol.* **6**, 34 (2010).

13. D. R. Swanson, *Perspect. Biol. Med.* **30**, 7 (1986).
14. D. R. Swanson, *Bull. Med. Libr. Assoc.* **78**, 29 (1990).
15. M. Weeber *et al.*, *J. Am. Med. Inform. Assoc.* **10**, 252 (2003).
16. D. R. Swanson, N. R. Smalheiser, *Artif. Intell.* **91**, 183 (1997).
17. M. Cokolr *et al.*, *Nat. Biotechnol.* **23**, 1243 (2005).
18. Z. S. Harris, in *Sublanguage: Studies of Language in Restricted Semantic Domains*, R. Kittredge, J. Lehrberger, Eds. (de Gruyter, Berlin, 1982), pp. 231–236.
19. Z. S. Harris, *The Form of Information in Science: Analysis of an Immunology Sublanguage* (Kluwer Academic, Dordrecht, Netherlands, 1989).
20. The National Library of Medicine’s Metathesaurus MetaMap connects more than 100 distinct medical terminologies.
21. J. D. Wren, *BMC Bioinform.* **5**, 145 (2004).
22. J. Feret, V. Danos, J. Krivine, R. Harmer, W. Fontana, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 6453 (2009).
23. M. Schmidt, H. Lipson, *Science* **324**, 81 (2009).
24. J. Bongard, H. Lipson, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 9943 (2007).
25. M. Twain, *A Connecticut Yankee in King Arthur’s Court* (Webster, New York, 1889).

Supporting Online Material

www.sciencemag.org/cgi/content/full/329/5990/399/DC1
Fig. S1

10.1126/science.1189416

PLANETARY SCIENCE

Carbonates and Martian Climate

Ralph P. Harvey

The case for global climate change on Mars is neither ambiguous nor arguable. Cold, dry, and dusty right now, the surface of Mars shows ancient fluvial valley networks, catastrophic flood channels, sediment-filled basins, and many other features that require stable liquid water at the planet’s surface. Such observations from orbit (1) imply that ancient Mars was substantially warmer than it is today. The favored scenario to heat up a young Mars is greenhouse warming driven by a thick, primitive CO₂-rich atmosphere of the kind thought common to all the terrestrial planets (2). Clearly, the greenhouse phase didn’t last long for Mars. As a result, almost every model for martian climate history includes an appreciable drawdown of CO₂ by precipitation of carbonates. However, only a small fraction of the carbonates predicted by this theory have been detected. On page 421 of this issue, Morris *et al.* (3) report the detection of ample carbonates dispersed within the subsurface, thereby strengthening arguments that CO₂ played a major role in the early martian climate system.

Carbon’s geochemical cycle, as originally modeled for Earth, involves interactions between the atmosphere and the hydrosphere that can dramatically influence global climate. Atmospheric CO₂ goes into solution in liquid water relatively easily and can stay there in cool conditions. This alone can draw down a large amount of CO₂. Furthermore, the resulting fluid, carbonic acid (H₂CO₃), reacts with the silicate minerals of the crust to weather the rock and release calcium, magnesium, and iron ions into the water. These ions can promote the precipitation of solid carbonates, (Ca,Mg,Fe)CO₃. Because carbonates remain stable under a great variety of near-surface conditions within the crust, they offer an effective and convenient storage medium for sequestering CO₂.

With the evidence for fluvial activity on early (and only early) Mars, progressive “fixing” of atmospheric CO₂ has been consistently favored by climate modelers. This theory adequately explains why, for the first billion years or so, Mars was warmed by an active greenhouse strong enough to support abundant liquid water and an active hydrologic cycle. Given our understanding of the links among atmospheric CO₂ levels, carbon sequestration, and greenhouse warming on Earth, the CO₂ drawdown explanation for martian cli-

The detection of large volumes of carbonates reasserts the importance of carbon dioxide in martian climate history.

mate change seems rational. Unfortunately, this martian climate jigsaw puzzle is missing a key piece: The widespread and/or massive deposits of carbonate minerals are nowhere to be found. Earth-based and orbital spectroscopic studies suggest at best an upper limit of a few percent within observable surface materials, around 20 to 30% of the amount needed for the accepted climate models (4, 5).

The result has been a cottage industry in explanations as to why we don’t see massive carbonate on Mars (6, 7). These explanations have two prominent themes. First is that our base-level assumption of an early martian climate primarily controlled by atmospheric CO₂ is too simplistic, either because other important chemical species were ignored or because the atmosphere simply never had the required mass. Certainly the former is true. Landers, rovers, and orbiters have directly detected an abundance of sulfate-dominated deposits, offering clear evidence that sulfur species played a role in the early martian climate (8). Like CO₂, atmospheric SO₂ is a greenhouse gas whose complex geochemical cycle includes dissolution into surface waters, acid weathering of surface rocks, and precipitation as relatively stable solids (9). CO₂ and SO₂ differ in stability and solubility under changing conditions of pH, pressure, and tem-

Department of Geological Sciences, Case Western Reserve University, Cleveland, OH 44106, USA. E-mail: rph@case.edu