ARTICLE IN PRESS

Cognition xxx (2014) xxx–xxx

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

Manifesto for a new (computational) cognitive revolution

Thomas L. Griffiths *

Department of Psychology, University of California, Berkeley, United States

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Computational modeling
Big data
Crowdsourcing

ABSTRACT

The cognitive revolution offered an alternative to merely analyzing human behavior, using the notion of computation to rigorously express hypotheses about the mind. Computation also gives us new tools for testing these hypotheses – large behavioral databases generated by human interactions with computers and with one another. This kind of data is typically analyzed by computer scientists, who focus on predicting people's behavior based on their history. A new cognitive revolution is needed, demonstrating the value of minds as intervening variables in these analyses and using the results to evaluate models of human cognition.

Over 60 years ago, the cognitive revolution made legitimate the scientific study of the mind (Gardner, 1987; Miller, 2003). Formal models of cognition made it possible to postulate processes that lie between a person's history and their actions, offering an alternative to the rigid stimulus-response structure of Behaviorism. Using new mathematical ideas – in particular, the notion of computation – a generation of researchers discovered a way to rigorously state hypotheses about how human minds work. I believe that we stand on the brink of a new revolution, with equally far-reaching consequences and an equally important role for computation. A revolution in how we *test* those hypotheses.

While the decades since the cognitive revolution have seen significant innovations in the kinds of computational models researchers have explored, the methods used to evaluate those models have remained fundamentally the same. In fact, those methods have arguably remained the same for over a century, being based on the small-scale laboratory science that characterized the first psychological research (Mandler, 2007). If you want to answer a question about the human mind (or publish a paper in *Cognition*) you formulate some hypotheses, bring an appropriate number of people into the laboratory, and have them carry out a task that distinguishes between those hypotheses.

But while we have remained focused on the events in our laboratories, the world outside those laboratories has changed. The internet offers a way to reach thousands of people in seconds. Human lives are lived more and more through our computers and our mobile phones. And the people with the most data about human behavior are no longer psychologists. They are computer scientists.

The mouse clicks and keystrokes of our online interactions are data, and figuring out how to make the best use of those data has become an important part of computer science. Recommendation systems that tell you which books you might be interested in, services that suggest related news stories, search engines that make use of the tags people apply to images, algorithms that select the advertisements you are most likely to click on... all are significant areas of research in computer science, and all are fundamentally based on the study of human behavior.

They are also all missed opportunities for cognitive science.

Recommendation systems need to divine human preferences – a problem that has been studied by both psy-

* Address: Department of Psychology, University of California, Berkeley, 3210 Tolman Hall # 1650, Berkeley, CA 94720-1650, United States. Tel.: +1 (510) 642 7134; fax: +1 (510) 642 5293.

E-mail address: tom_griffiths@berkeley.edu

http://dx.doi.org/10.1016/j.cognition.2014.11.026
0010-0277/

Please cite this article in press as: Griffiths, T. L. Manifesto for a new (computational) cognitive revolution. *Cognition* (2014), http://dx.doi.org/10.1016/j.cognition.2014.11.026

chologists and economists (Lucas et al., 2014). Identifying related news stories requires extracting appropriate representations of the meaning of text, a key problem in studying language and memory (e.g., Landauer & Dumais, 1997; Jones & Mewhort, 2007). Image tagging is a problem of categorization, a central topic in cognitive psychology (e.g., Rosch, 1978; Medin & Schaffer, 1978; Nosofsky, 1986). And predicting what advertisements people will click on involves combining preferences, semantic representations, and categorization – something that would seem to require a rich model of human cognition.

Except that is not how computer scientists solve these problems. In practice, recommendation systems are typically based on "collaborative filtering" – predicting what you will purchase based purely on the similarity of your behavior to the behavior of others, not on building a complex model of your preferences (e.g., Linden, Smith, & York, 2003). Systems for processing text and images are evaluated via their performance on information-retrieval tasks (such as how often they identify a document or image somebody might be searching for), rather than being compared against richer metrics based on human cognition (this difference can be seen for a single model in Blei, Ng, & Jordan, 2003; Griffiths, Steyvers, & Tenenbaum, 2007). And the advertisements you see on webpages are chosen by reinforcement-learning algorithms that infer what people are likely to click on based on the webpages they recently visited and the content of the current webpage (e.g., Pandey & Olston, 2006).

All too often, behavioral data is analyzed as just that – *behavior*. And, as a result, the theoretical assumptions underlying these analyses would not seem controversial to a Behaviorist: that people act similarly to one another, and that future actions can be predicted from past actions.

Hence this call to revolution. This call to a new *cognitive* revolution. To take back behavioral data, and – just as in the last cognitive revolution – to demonstrate the value of postulating a mind between browsing history and mouse movements.

Ubiquitous records of human behavior offer the potential to study human cognition at a scale and level of validity that could never be achieved in the laboratory. To take just one example, Yahoo! recently made available (at http://labs.yahoo.com/news/yfcc100m/) 100 million images together with the tags that had been applied to those images by users – more data than has ever been collected in laboratory studies of categorization, using real images rather than artificial stimuli. Services like Twitter offer access to the stream of consciousness of millions of people, while Facebook provides information about their connections and interactions. Location trackers in mobile phones reveal where we go, and motion trackers reveal what we do when we get there.

My hope is that cognitive scientists can use this kind of data not just to get insight into how human minds work, but to improve the strategies that computer scientists have developed for working with these datasets – to leverage our decades of experience in thinking about the causes of human behavior to develop richer, more cognitive models that lead to better predictions.

There are already lines of research that have begun to explore the potential of these new sources of data about the mind. First of all, my characterization of the methods of modern psychology is a few years out of date – increasingly, psychologists are making use of crowdsourcing services such as Amazon's Mechanical Turk to run experiments over the internet at a larger scale than would be possible in the laboratory (Crump, McDonnell, & Gureckis, 2013; Buhrmester, Kwang, & Gosling, 2011; Mason & Suri, 2012). Researchers have begun to use large databases of naturalistic images in psychological experiments, offering strong tests of psychological theories (e.g., Isola, Xiao, Torralba, & Oliva, 2011; Abbott, Austerweil, & Griffiths, 2012). Others have explored the question of how human categorization could be studied using online databases (Glushko, Maglio, Matlock, & Barsalou, 2008). Computer games – with hundreds of thousands of players – offer a different way to study skill acquisition (Stafford & Dewar, 2014). And records of financial transactions have begun to be used to inform theories of economic decision-making (Stewart, Chater, & Brown, 2006; Gelman, Kariv, Shapiro, Silverman, & Tadelis, 2014).

This new revolution will face challenges. Crowdsourcing of experiments is potentially transformative, offering a way to make progress in studying the mind at a speed and level of precision that has not previously been possible. But it is also an under-exploited resource. We need to stop viewing crowdsourcing as a way to do what we used to do in the laboratory more quickly and at a larger scale, and start thinking about how it changes what we *can* do. For the first time, researchers using behavioral methods to study the mind have a tool that has the same high-bandwidth, high-cost profile as neuroimaging: we could spend a few hundred dollars for an hour on an MRI machine, but we might get richer and more meaningful data by spending the same amount on Mechanical Turk. And we should write grant proposals that give this kind of intensive behavioral data collection equally high priority.

More importantly, knowing that you can easily have thousands of participants in a study should change how studies are designed. Rather than simply scaling up an experiment intended to provide a single bit of information – which of two hypotheses is correct – we need to develop new experimental paradigms that give us a richer picture of human cognition. For example, my collaborators and I have adapted algorithms that computer scientists and statisticians use for sampling from complex probability distributions to define new experimental methods that can be used to estimate distributions associated with human category representations (Sanborn, Griffiths, & Shiffrin, 2010) and prior distributions (Lewandowsky, Griffiths, & Kalish, 2009; Yeung & Griffiths, 2011; Canini, Griffiths, Vanpaemel, & Kalish, 2014). These methods require large numbers of participants (or many judgments per participant), but provide a great deal of insight into the mental representations that inform people's judgments.

Exploring the research potential of large-scale behavioral datasets is also challenging. While these datasets offer a depth and realism that goes far beyond that of laboratory data, they do so at the cost of making it harder to identify causality. In working with these datasets, we need to adopt a different mindset – more like the mindset of an astronomer, making sense of noisy data viewed from far away.

This does not mean that the data are not scientifically useful – the success of astrophysics bears this out. But it does mean that we should combine the old with the new, using laboratory (or online) experiments to establish causality in cases where it is not clear from existing records of behavior.

Pushing this analogy with astrophysics further, it is instructive to consider how astronomical observations and laboratory experiments are combined to make scientific progress: through theory. Theoretical physics provides the broader context that links small-scale laboratory science with large-scale observations, allowing hypotheses to be evaluated using both of these methods. In the same way, theories of cognition that make predictions at multiple scales – in the lab, and in the wider world – are going to be fundamental to integrating these new sources of data into the practice of cognitive science.

Computation thus plays an intimate role in this new revolution. As in the last revolution, it is what allows us to formulate precise hypotheses about minds and their consequences for behavior. But the six decades that have passed since that last revolution have brought the theoretical construct of computation into an applied reality, and that reality means computers are also the medium by which the data we use to evaluate our hypotheses will be collected. They are also the medium by which the data will be analyzed, moving into a more central role in psychological research. If you want to train a student to be in a position to make the most of this new revolution, programming skills will play as important a role in the curriculum as experiment design.

Just as the current state of cognitive science would have been hard to predict in the 1950s, it is hard to predict what the consequences of this new cognitive revolution might be. But my vision is of a very different kind of laboratory for studying the mind – one where the rich information about our behavior that is collected ubiquitously by our devices is fodder for evaluating theories of cognition, and where theories of cognition play a central role in how that information is used. Over the next few decades, I hope that the papers that appear in *Cognition* will increasingly make use of data collected outside our laboratories, and that the theories they present will increasingly have an impact that goes far beyond cognitive science.

Acknowledgements

Preparation of this paper was supported in part by Grant No. SMA-1228541 from the United States National Science Foundation. I thank Nick Chater, Tania Lombrozo, and the participants in the Data on the Mind workshop for helping to shape these ideas.

References

Abbott, J. T., Austerweil, J. L., Griffiths, T. L. (2012). Constructing a hypothesis space from the web for large-scale Bayesian word learning. In *Proceedings of the 34th annual conference of the cognitive science society*.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research, 3*, 993–1022.

Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science, 6*, 3–5.

Canini, K. R., Griffiths, T. L., Vanpaemel, W., & Kalish, M. L. (2014). Revealing human inductive biases for category learning by simulating cultural transmission. *Psychonomic Bulletin and Review, 21*, 785–793.

Crump, M. J., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PloS One, 8*, e57410.

Gardner, H. (1987). *The mind's new science: A history of the cognitive revolution*. New York: Basic Books.

Gelman, M., Kariv, S., Shapiro, M. D., Silverman, D., & Tadelis, S. (2014). Harnessing naturally occurring data to measure the response of spending to income. *Science, 345*, 212–215.

Glushko, R. J., Maglio, P. P., Matlock, T., & Barsalou, L. W. (2008). Categorization in the wild. *Trends in Cognitive Sciences, 12*, 129–135.

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic association. *Psychological Review, 114*, 211–244.

Isola, P., Xiao, J., Torralba, A., Oliva, A. (2011). What makes an image memorable? In *IEEE conference on computer vision and pattern recognition (CVPR)* (pp. 145–152).

Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review, 114*, 1–37.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review, 104*, 211–240.

Lewandowsky, S., Griffiths, T. L., & Kalish, M. L. (2009). The wisdom of individuals: Exploring people's knowledge about everyday events using iterated learning. *Cognitive Science, 33*, 969–998.

Linden, G., Smith, B., & York, J. (2003). Amazon.com recommendations: Item-to-item collaborative filtering. *Internet Computing, 7*, 76–80.

Lucas, C. G., Griffiths, T. L., Xu, F., Fawcett, C., Gopnik, A., Kushnir, T., et al. (2014). The child as econometrician: A rational model of preference understanding in children. *PloS One, 9*, e92160.

Mandler, G. (2007). *A history of modern experimental psychology: From James and Wundt to cognitive science*. Cambridge: MAMIT Press.

Mason, W., & Suri, S. (2012). Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods, 44*, 1–23.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85*, 207–238.

Miller, G. A. (2003). The cognitive revolution: A historical perspective. *Trends in Cognitive Sciences, 7*, 141–144.

Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General, 115*, 39–57.

Pandey, S., Olston, C., (2006). Handling advertisements of unknown quality in search advertising. In *Advances in neural information processing systems* (Vol. 19, pp. 1065–1072).

Rosch, E. (1978). Principles of categorization. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.

Sanborn, A. N., Griffiths, T. L., & Shiffrin, R. (2010). Uncovering mental representations with Markov chain Monte Carlo. *Cognitive Psychology, 60*, 63–106.

Stafford, T., & Dewar, M. (2014). Tracing the trajectory of skill learning with a very large sample of online game players. *Psychological Science, 25*, 511–518.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology, 53*, 1–26.

Yeung, S., Griffiths, T. L. (2011). Estimating human priors on causal strength. In *Proceedings of the 33rd annual conference of the cognitive science society*.