TOPICS
TOPICS IN COGNITIVE SCIENCE

Topics in Cognitive Science 8 (2016) 548–568

ISSN:1756-8757 print / 1756-8765 online
DOI: 10.1111/tops.12212

This article is part of the topic "Discovering Psychological Principles by Mining Naturally Occurring Data Sets," Robert L. Goldstone and Gary Lupyan (Topics Editors). For a full listing of topic papers, see: http://onlinelibrary.wiley.com/doi/10.1111/tops.2016.8.issue-3/issuetoc.

Discovering Psychological Principles by Mining Naturally Occurring Data Sets

Robert L. Goldstone,[a] Gary Lupyan[b]

[a]*Department of Psychological and Brain Sciences, and Program in Cognitive Science, Indiana University*
[b]*Department of Psychology, University of Wisconsin-Madison*

Received 11 March 2016; received in revised form 20 March 2016; accepted 20 March 2016

Abstract

The very expertise with which psychologists wield their tools for achieving laboratory control may have had the unwelcome effect of blinding psychologists to the possibilities of discovering principles of behavior without conducting experiments. When creatively interrogated, a diverse range of large, real-world data sets provides powerful diagnostic tools for revealing principles of human judgment, perception, categorization, decision-making, language use, inference, problem solving, and representation. Examples of these data sets include patterns of website links, dictionaries, logs of group interactions, collections of images and image tags, text corpora, history of financial transactions, trends in twitter tag usage and propagation, patents, consumer product sales, performance in high-stakes sporting events, dialect maps, and scientific citations. The goal of this issue is to present some exemplary case studies of mining naturally existing data sets to reveal important principles and phenomena in cognitive science, and to discuss some of the underlying issues involved with conducting traditional experiments, analyses of naturally occurring data, computational modeling, and the synthesis of all three methods.

Keywords: Big data; Research; Language; Decision-making; Perception; Memory; Representation; Statistics

Correspondence should be sent to Robert Goldstone, Department of Psychological and Brain Sciences, 1101 E. 10th St., Indiana University, Bloomington, IN 47405. E-mail: rgoldsto@indiana.edu

1. Introduction

With more than 100 years of collective practice, experimental psychologists have become highly sophisticated in their application of well-controlled, laboratory experiments to reveal principles underlying human cognition and behavior. The resulting careful experimental designs, controls, counterbalancing, and proliferation of novel methods should be valued and encouraged. However, the very expertise with which psychologists wield their tools for achieving laboratory control may have had the unwelcome effect of blinding us to the possibilities of discovering principles of behavior without conducting experiments. One of the most promising areas of future growth in this direction is achieved by analyzing naturally occurring real-world data sets that affect and reveal human behavior.

Large and systematically collected sets of behaviorally relevant data have existed for many decades. These include census records, library classification systems, dictionaries and text collocations, financial decisions to save and invest, and Department of Labor statistics. In recent years, however, the diversity, availability, and scale of these data sets have exploded. Analysis of "Big Data" has already transfigured economics, sociology, biology, and other sciences, as indicated by the "Big Data" initiative of the United States government (http://blogs.sciencemag.org/sciencecareers/2012/03/new-federal-big.html). A risky inductive leap is not required to predict that cognitive science is also poised to be similarly affected by Big Data techniques (Jones, 2016). Although our emphasis is on analysis of naturally occurring data sets (NODS), there is considerable overlap between our topic and Big Data analysis because the size of useful, naturally occurring data sets has been growing impressively with time. When creatively interrogated, a diverse range of data sets provide powerful diagnostic tools for revealing principles of judgment, perception, categorization, decision-making, language use, inference, problem solving, and mental representations.

The immediate goal of this topic is to present some exemplary case studies of mining large data sets to reveal important psychological principles and phenomena. The broader goal is to stimulate cognitive scientists to consider novel ways of harnessing the power of large data sets for informing their own areas of inquiry into mental processes. Our understanding of minds can be greatly enhanced by increasing our awareness of the kinds of possible data sets and analyses available to us. Examples of these data sets include dictionaries, logs of group interactions, collections of images and image tags, text corpora, history of financial transactions, Wikipedia edit histories, co-occurrence of terms on web pages, trends in twitter tag usage and propagation, historical records, demographics, patent use and dependencies, jury and judge decisions, mobile phone calls, the reading and forwarding of online news stories, consumer product sales, performance in high-stakes sporting events, dialect maps, and scientific citations. Behavioral economists have been effectively plumbing field studies for several decades now (DellaVigna, 2009), and so it could be argued that cognitive scientists are coming late to an already well-picked-over banquet. However, the theoretical preoccupations of cognitive scientists and economists are sufficiently different that we expect that cognitive scientists may even be able to look at old, well-studied data sets with fresh eyes to see new things. For example, Berger and Pope (2011) examined 18,000 professional and 45,000 collegiate basketball games to find that teams behind one point at halftime are more likely to win than those ahead by one point. This result points to the powerful motivating effect of feeling that it is possible to overcome one's setbacks. Moreover, the strength of this motivation can be quantified—it leads to about 6% more victories than would otherwise be predicted. A myriad of patterns like this lie latent in data sets, needing only behavioral scientists endowed with inquiring minds and relevant theoretical perspectives to be brought to the surface.

In addition to presenting case studies of NODS for revealing principles underlying human behavior, the authors also consider fundamental issues regarding the use and analysis of large data sets. For example, one important topic is ways of statistically analyzing data when factors cannot be experimentally manipulated, but we must make do with the data we have. Although large real-world data sets are not as hygienically controlled as their laboratory counterparts, and do not allow for true random assignment and manipulation of factors, their sheer size can allow factors to be statistically pulled apart even when they cannot be manipulated. Particularly impressive are recent nonlinear state space reconstruction methods for deriving causality from correlated time series in complex ecosystems with many interacting species (Sugihara et al., 2012). For example, the technique has been applied to sardine and anchovy populations over time that show one species' increase is correlated with the other species' decrease. By probing the historical record to see whether several variables, including past sardine populations, predict future anchovy populations and vice versa, the researchers determined that the species are probably not directly influencing each others' numbers, but rather both are influenced by ocean surface temperature.

Methods like this show that causal relations across variables can be deduced by patterns of covariation over time (cf. Granger, 1969), and in many cases, having many recorded variables makes it easier, not harder, to determine causal relations among the interacting variables (Pearl, 1988). While conventional wisdom maintains that "You can't derive causation from correlation," modern statistical and machine learning analysis methods show that one can probabilistically determine causal relations from correlations (more precisely, conditional dependencies) if one has many interrelated variables and one makes some plausible assumptions about how causal processes operate. These analyses have made striking progress toward understanding causal relations in medical diagnosis, homeland security screening, automated user assistance, genetic counseling, natural language understanding, and mapping gene expression data (see Pearl, 2009 for a review). The cognitive scientist wishing to conduct research with NODS can benefit from developments in many analysis techniques: correlation, regression, multidimensional scaling, clustering, network analysis, kernel density estimation, Markov models, autoregressive models, detrended fluctuation analysis, Directed Acyclic Graphs, Granger causality, Bayes nets, and machine learning algorithms. For most theoretically motivated questions, there are appropriate and readily available analysis techniques.

2. NODS, experiments, and models

Despite some of the advantages of NODS we just highlighted, large data sets with many variables do not always compensate for stimulus control and ability to manipulate variables that characterize genuine laboratory experiments. The psychophysicist interested in color perception would be foolish to wait around for the natural world to provide a set of scenes with exactly the right color comparisons and contrasts to test her theory. There are undeniable benefits derived from being able to experimentally *intervene* on a system by systematically turning on/off some variables while leaving all others intact (Pearl, 2000). One of the themes that emerges from many of the articles in this issue is that NODS should *supplement*, not *supplant*, experiments. The articles describe best methods of practice for determining whether a question about minds is better answered using controlled laboratory experiments or analyses of large data sets (Pope, 2016). Some of the primary uses of NODS for cognitive scientists interested in how individual minds work are shown in Table 1. One use of NODS is to *provide external validation* for experiments that have been primarily tested in the laboratory. For example, Berger (2016) tests whether primacy effects observed in the laboratory (e.g., better memory for words occurring earlier in a list) are found in a domain that academicians care deeply about—citations to articles appearing in scholarly journals. He finds, in fact, that articles appearing earlier in a given issue of a journal are more likely to be cited than articles appearing later (hence the placement of our introduction in this issue of *topiCS*). Moat, Olivola, Chater, and Preis (2016) describe another example of external validation of a theory. According to the "decision making by sampling" theory, people make probability judgments by sampling events in their memories that belong and do not belong to a target category. Consistent with the theory, Moat et al. find that people overestimate the probability of dying in a way that kills many people at once, because these events tend to be overrepresented in news media, including the Google News Archive. One valuable role of NODS for external validity is that their use can show whether or not a laboratory observed result has a sufficiently large effect size so as to influence actual behaviors of people making consequential real-world decisions. It is one thing to show in a laboratory that people tend to simplify their lives and reduce their cognitive work by using the default choices given to them. It is another thing to show that they use these default options even when a making decision about their own retirement savings that will affect by hundreds of thousands of dollars how much money they have when they retire

Table 1
Uses of naturally occurring data sets for research

External validation of laboratory experiments
Demonstrate phenomena that motivate follow-on laboratory research
Discover patterns of information latent in environments
Create stimuli for experiments
Construct and test computational models of cognition

(Madrian & Shea, 2001; Pope, 2016). Tests of psychological theories using NODS allow researchers to tell whether their ideas have real-world, not only statistical, significance.

Another use of NODS is to *motivate subsequent laboratory research to distinguish theoretical accounts of real-world outcomes*. Many university researchers assume that the natural sequence of research is one of "scaling up." One finds a result in the laboratory and after one is convinced that it can be reliably found in the lab, it is scaled up to see if it can be found in the world at large. This progression is one of the standard models for how Department of Education projects are expected to evolve—from basic research on learning, to eventual classroom application. However, an equally important research enterprise is "scaling down"—finding a natural phenomena and taking it into the laboratory to determine what makes it work (Nathan & Alibali, 2010; Nathan & Sawyer, 2014). For example, Christiansen and Monaghan (2016) find subtle differences in the morphological and phonological properties for nouns versus verbs in dictionaries and corpora of child-directed speech. This leads naturally to the question, "Are people sensitive to these phonological differences?" They report follow-up laboratory experiments that give a positive answer to this question, showing that even with relatively little exposure to a language, children become sensitive to phonological cues that distinguish nouns from verbs and can use them to infer the lexical status of novel words. As another example, Pope (2016) reports a striking example of a left-digit bias in used car purchases, whereby buyers pay far more for a car with only 9,970 miles compared to one that has 10,020 miles—a much bigger difference than would be predicted given the average buying price difference between, say, cars with 9,930 versus 9,980 miles. Follow-up laboratory experiments were then constructed to distinguish between different cognitive mechanisms underlying this effect. Through these experiments the researchers end up supporting the hypothesis that the left-digit bias is driven by inattention because participants' recall for non-left digits of a car's price was far worse than for left digits. One of the advantageous aspects of NODS for motivating subsequent laboratory research is that the researchers can be confident that the effect they experimentally track down is strong enough to drive real-world patterns because it is those patterns that inspired the laboratory investigation in the first place.

A third use of NODS is to discover patterns of information latent in environments. This use of NODS is surprisingly reminiscent of ecological psychology (Gibson, 1979; Neisser, 1976), though with markedly different methods and considerations concerning what counts as a person's surrounding ecology. Consistent with studying NODS, the ecological psychology movement has always emphasized studies of real-world behavior as opposed to the artificial environment of the laboratory. J. J. Gibson stressed the importance of studying the environment in which behavior takes place, and the often hidden or subtle perceptual cues that an organism could use to behave in an adaptive manner. It was a surprise to us to see how well that description fits much of the current use of NODS as represented by the articles in this issue. So much so, in fact, that we are prompted to coin the term "cultural neo-ecological psychology" (CNP) to describe the use of NODS to quantify the informational ecology that surrounds human behavior. This ecology is increasingly becoming shaped by cultural products. One example of this approach is Christiansen and Monaghan's (2016) quantitative articulation of the phonological cues that predict part of speech in English. A second example is Heit and Nicholson's (2016) study of the positions and attributes of Democrats and Republicans. A third example is Griffiths, Abbott, and Hsu's (2016) application of machine learning techniques to reveal how natural categories are structured. As a fourth example, Moat et al. (2016) study the distribution of web pages that are prominently available to people as they forage their web environment for information. As a fifth example, Vincent-Lamarre et al. (2016) analyze information about connections between words and the minimal networks of words sufficient to define all other words, as revealed by large corpora analyses of dictionaries.

In at least the cases of Christiansen, Heit, Griffiths, and Moat, the researchers are explicitly interested in showing that people are sensitive to particular sources of distributional information present in the real world. These researchers are precisely interested in whether and how people internalize the distributional information available to them, and so depart from the traditional ecological psychology exhortation to "Ask not what's inside your head, but what your head's inside of" (Mace, 1977). This focus was presaged by Shepard's (1984) theorizing that the constraints of internal representations echo those of the cognizer's ecological context. Another obvious difference in emphasis between CNPs and traditional ecological psychologists is that the former are focusing on the structure of an external world that is largely based on information on the web or social constructs—words (Christiansen & Monaghan; Vincent-Lamarre), prices (Pope), stocks (Moat et al.), mathematics problems (Koedinger, Yudelson, & Pavlik, 2016), and scholarly articles (Berger). As more of our attended world migrates online, it makes sense that the structure of this online world becomes of central psychological concern. Ranked search results are just as much part of the modern web citizens' experience as is the optic flow studied by Gibson.

A fourth use of NODS is to create stimuli for experiments. This use most directly belies the view of NODS and lab experiments as existing in a competitive or zero-sum relation. Several of the articles in this issue use NODS in order to create laboratory stimuli with an aim toward making the experiments more externally valid. Heit and Nicholson (2016) asked survey respondents to estimate the probability that a candidate with a particular profile description was a Democrat or Republican. The descriptions were generated by a large survey from the American National Election Study that identified people's political parties, demographic attributes, and beliefs. In this manner, the researchers were able to ascertain that people are more accurate in their ability to identify Democrats and Republicans than would be expected by some "uninformed voter" accounts. That is, people are sensitive to actual differences in positions and attributes between Democrats and Republicans. An analogous logic was employed by Christiansen and Monaghan (2016) to show that children are sensitive to actual phonological differences between nouns and verbs. Brady, Konkle, Alvarez, and Oliva (2008) sampled stimuli in a long-term memory experiment using results from Google Image Search queries in order to provide experiment participants with stimuli that have natural statistical distributions and levels of variation (see Griffiths et al., 2016). If the use of NODS for human experiments is rapidly increasing, it is fair to say that for machine learning experiments it is absolutely exploding. The tremendous recent advances in deep learning depend crucially on giving machine algorithms large and representative data sets of naturally occurring image, audio, and text samples. The ability of the algorithms to derive latent statistical structures at multiple levels embedded within data is truly impressive (Bengio, Courville, & Vincent, 2013; LeCun, Bengio, & Hinton, 2015), and so effective that it would be surprising if humans (and animals more broadly) had not evolved to exploit similar methods.

Related to this confluence between human and machine learning, a fifth use of NODS is to inform the construction not of experiments, but of computational models of cognition. Koedinger et al. (2016) show how the choice between different cognitive models of mathematical reasoning can be adjudicated with the help of large amounts of data, courtesy of the Datashop repository of educational technology data (Koedinger, Stamper, Leber, & Skogsholm, 2013). The particular data they studied concerned the patterns of successes and errors that students make across a suite of related mathematical problems. The models that they considered varied in how general their knowledge components were, and the empirically favored models fell between the extremes of positing a single faculty of "general intelligence" and positing a different knowledge element for every task. Models that posited knowledge components at an intermediate cognitive level were most strongly supported, and the data were able to disconfirm some natural alternative hypotheses about what these components might be. Griffiths et al. (2016) and Vincent-Lamarre et al. (2016) also use NODS to distinguish between alternative computational accounts of cognition.

The articles in this issue of *topiCS* describe converging methods to validate experimental results with real-world data sets and refine theories generated by analyzing real-world data sets by conducting controlled experiments. None of the authors suggested that experiments should be superseded by NODS, but rather that there will be a bidirectional process in which each activity informs and constrains the other. One final hybrid strategy for bridging experiments and NODS analysis that was not represented among the articles but is likely to become increasingly important for cognitive science is to create real and virtual world services and environments for users that allow a researcher to observe users' natural behavior yet still explicitly manipulate key variables that could affect users' behavior (Salganik, Dodds, & Watts, 2006; Szell, Lambiotte, & Thurner, 2010). Salient examples of this approach have been offered by Facebook's research arm, which has shown that it can affect the moods of Facebook citizens by varying the positive versus negative content of posts visible to them on their news feeds (Kramer, Guillory, & Hancock, 2014), and it can affect actual voting behavior of facebook citizens that they target with political mobilization messages *and* their networked friends (Bond et al., 2012). Synthetic worlds offer particularly exciting opportunities for manipulating social variables in vivo *virtualis* because they allow us to experiment with alternative foundational organizational choices that would be prohibitively costly and infeasible to implement otherwise (Castranova, 2005).

3. Advantages of NODS

Although some of the advantages underlying NODS have already been implicitly described in the section above, we wish to explicitly enumerate several of the reasons why NODS are particularly exciting at the present moment for cognitive science:

A. *Big Data technical innovations*. Cognitive scientists stand to benefit by taking advantage of the many methodological tools now readily available for analyzing Big Data. Technological innovations like MapReduce and Hadoop make processing large data sets tractable through parallel architectures. Otherwise unwieldy data sets like a month's worth of twitter feeds, the network of citations among all journal papers, and a billion words of freely available Google books can now be mined by social scientists without needing a PhD in computer science.
B. *Wide availability of interesting data sets*. There has been an exponential growth in the number and size of freely available data sets. Some prominent data sets include scholarly citations, telephone calls, the movement of currency, disease spread, gossip spread, patterns of collaboration, patent uses and dependencies, jury decisions, business transactions, census records, and raw data from national surveys. Table 2 is provided in the spirit of an "inspiration pump" to identify for cognitive scientists some fertile data sets for particular fields, although the list only begins to scratch the surface of possibilities. Beyond these online data sets, there are increasingly many sensors tracking our movements and behaviors, and these have provided a wealth of data about habits, motivation, social networks, and factors affecting performance. Many of the specific data sets interrogated by the topic's authors did not exist 10 years ago, nor did the technology required to assemble the data sets. This is notably true for the researchers using Google Ngrams, GWAP, Flickr, and Twitter tags and follows, to take a few examples. Moreover, paralleling the exponential increase in the number of open data sets, there has been a sharp increase in the tools available to analyze the data sets. Tools like Beautiful Soup greatly expedite the collection of web data, as do open source bots to crawl data, easily operated database management and query systems, and advances in natural language processing.
C. *Data on socially and personally important behaviors and decisions*. The human decisions revealed by NODS are often highly consequential ones and are made by highly motivated individuals. Decisions like how to wager on a game show (Post, Assem, Baltussen, & Thaler, 2008), where to aim one's golf ball in a professional tournament (Pope & Schweitzer, 2011), what to name one's child (Berger, Bradlow, Braunstein, & Zhang, 2012; Berger & Le Mens, 2009; Gureckis & Goldstone, 2009), or how to save or invest one's money (Bernartzi & Thaler, 2007) are major life-changing decisions that people take, or at least ought to take, very seriously. These data involve decisions and behaviors that have direct and often times profound ramifications for both individuals and societies.
D. *Avoiding experiment-related contamination*. NODS typically allow researchers to investigate natural behavior in ways that are uncontaminated by perceived task demands, presentation biases, and researcher expectancies. Admittedly, NODS often have the converse *dis*advantage—reflecting artifacts due to specific "nuisance" factors in the real-world situation. One example of this was cited by Pope (2016)—buyers of cars on the used car auction market may not intrinsically value cars with 9,990 miles much more than cars with 10,010 miles, but instead believe that *their* customers will, and so are willing to pay appreciably more for them. Koedinger et al. (2016) describe

Table 2
A small subset of available data sets with potential relevance for cognitive science. Citations and Internet links have been omitted because an Internet search will readily retrieve up-to-date links

Fields and Data Sets	Sample Questions
Language: dictionaries, thesauri, text corpora, Wordnet, dialect maps, Google Ngrams, Google Word2Vec; WALS, OpenSubtitles, CYC, CHILDES	How does the composition of speakers of a language affect the evolution of the language? How does status differential, exposure, and familiarity influence a person's dialect?
Search, imitation, and exploration: Social Security records of baby names, journal citations, patents, Twitter tags, web memes, web links, fishing vessel routes, music downloads, article views, children's art within and between classrooms	How are people affected by similarity, popularity, momentum, and success when deciding what memes to imitate? What heuristics guide people's decisions to download an article or piece of music?
Cognition: tutoring system interactions, standardized tests, cognitive training games	What are the separable knowledge components that comprise proficient mathematical reasoning? What does the relation between two tasks have to be for training in one task to lead to improvements in the other?
Vision and object recognition: Flickr, GWAP, natural image statistics over photographs, ImageNet, browser click streams, movie trailers	What kinds of hierarchical visual representations arise when learning systems are fed natural scenes with realistic distributions? Are human object representations veridical representations or warped to emphasize diagnostic visual features?
Decision making: retirement fund holdings, financial transactions, sales records, Ebay auction bids, NFL trades, credit card balances, donations	How influenced are people by their immediate context when making long-term decisions? Does the value of a player within their original, longtime team predict their value when traded to another team? Why do people put a greater value on an object valued by others?
Performance: golf scores, chess ratings, race times, home and away games, time course of points in a game	If a person wishes to maximally improve their performance, what is the optimal relative skill level of their opponent? What is the optimal level of stress for performance?
Group behavior and social networks: Wikipedia edits, currency movements, cell phone calls, paper collaborations, United Nations resolutions, General Social Survey	What causes some discussions to split into antagonistic sides and others to reach consensus? How does geographic and semantic distance influence the likelihood and profitability of a social interaction?

(continued)

Table 2 (continued)

Fields and Data Sets	Sample Questions
Physiological measurements: emotion detection from video, DNA, walking speed, activity monitors, heart rate monitors, blood chemistry, health club attendance	What are the genetic components of decisions to explore or exploit? What factors motivate people to start and keep exercising?

another nuisance factor in their NODS; because they use an adaptive mastery algorithm, the number of students contributing to an error rate estimate decreases as the number of practice opportunities increases. Braithwaite, Goldstone, van der Maas, and Landy (2016) describe a related measurement problem with educational systems that adapt the problems given to students based on the students' performance. Whether the situational artifacts are worse than the laboratory artifacts will depend on the strength of task demands. For situations in which people are motivated to present themselves in a favorable light, choke when they know they are being observed (Beilock & Gray, 2007), or make decisions differently when explicitly reflecting on them versus acting intuitively, the benefits of employing NODS will often more than outweigh their costs.

E. *Understanding cross-level interactions between individuals and the social structures in which they are embedded*. NODS allow us to investigate the bidirectional interplay between internal psychological processes and the external artifacts produced by those processes, including language structures (Griffiths, Steyvers, & Tenenbaum, 2007; Monaghan, Christiansen, & Fitneva, 2011; Steyvers & Tenenbaum, 2005), laws, political structures (Adamic & Glance, 2005), group organizations (Cooke & Hilton, 2015), and creative works. The study of NODS fits well with an "extended mind" (Clark, 2004, 2009; Hutchins, 1995) approach to cognitive science that takes seriously the bidirectional links between individual human beings and their surrounding environments for establishing distributed cognitive units. If cognition emerges as a function of the interactions between individual humans and their environments, then NODS play an indispensable role in determining the environmental components that shape these distributed cognitive processes.

4. Five exemplary case studies

In addition to the articles in this issue, we would like to summarize five case studies of the use of NODS, both as inspirations for their possible theoretical diagnosticity and also because of the generalizable principles they suggest. Although the selection of only five examples is necessarily arbitrary and hopelessly incomplete, these case studies are historically and societally significant.

4.1. Anderson and Schooler (1991)

Anderson and Schooler present an early example of NODS by analyzing New York Times headlines, the CHILDES database of utterances spoken to and by children, and 4 years' worth of email to John Anderson. They enlisted these data sets to explain why our memories may show the particular pattern of forgetting that they do. On the left panel of Fig. 1, they show an approximate power law relation between the probability of a word appearing in a New York Times headline and the number of days since the word last appeared in the New York Times. They show a similarly shaped power law relation between John Anderson's probability of receiving an email from a particular person and the number of days since last receiving an email from that person. This plot shows the "burstiness" of our news. If the news covers a particular topic on a day, then it is likely to cover that topic in the near future. If we have recently received an email from somebody, then we are "at risk" for receiving still more emails from them in the near future. Often times we will actively email a person many times within a short period of time only to let the correspondence lie fallow for many months until the next burst of activity. The right panel of Fig. 1 shows Ebbinghaus's (1885) power law relation between the retention of a word measured by savings and the number of hours since the word was last studied. On the basis of close correspondences like this, Anderson and Schooler argue that our memory systems may have evolved so that they are well matched to natural environmental

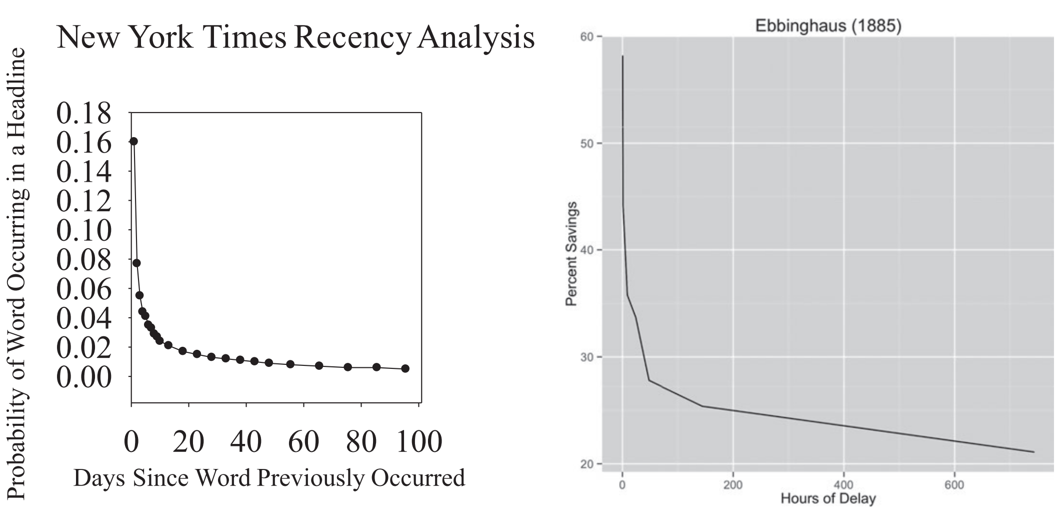

Fig. 1. Data analyzed by Anderson and Schooler (1991, 2000). The left panel shows the approximate power law relation between the probability of a word appearing in a New York Times headline and the number of days since the word last occurred. The right panel shows the power law relation observed by Ebbinghaus (1885) between the retention of a word measured in savings and the time since the word last studied. Figure reprinted with permission from the authors.

statistics. Our memory for something deteriorates in almost exact proportion with the likelihood of needing that memory.

4.2. Michel et al. (2011)

In 2011, the Google Books team introduced their transformative Ngram Viewer, which allows anybody to search for phrases in more than 5 million books and half a trillion words. This resource has a large impact on research in linguistics, economics, political science, law, sociology, and psychology. One of the theories that they tested was that the past tense forms of verbs should regularize over time, with the –ed form ascending in dominance over time. The results shown in Fig. 2 do not generally support this theory. Michel et al. observed far more idiosyncratic trajectories for verbs than would be expected from a general irregular-to-regular shift over time. Some word (shown in red) became significantly more irregular over time. “Lighted” and “Waked” sound strange to our present day ears but would not have sounded strange in 1850. Although there is not a cross-dictionary trend toward regularization, there are distinct local pockets of regularization. For example, of the six words that shifted from >50% irregular to >50% regular (shown in blue), four of them belong to the “-t” cluster that creates past tenses like bend->bent, build->built, send->sent, spell->spelt, and smell->smelt. Other analyses from this same set of verbs show that there is a relation between regularization over time and verb frequency, with a tendency for lower frequency verbs to regularize over time compared with high-frequency verbs. One account for this is that it is hard to remember an irregular form for a verb if that irregular form is seen only rarely. Exceptions to the standard “-ed” are more likely to evade regularization if they maintain a fairly high frequency (though see Lupyan & McClelland, 2003 for an alternative account).

4.3. Camerer, Babcock, Loewenstein, and Thaler (1997)

Camerer et al. studied the driving logs of New York City taxi drivers across many days, finding that as taxi drivers earn higher hourly wages in a day, they also tend to work fewer hours. This is shown by the negative correlations in Fig. 3 between hourly wages and number of hours worked. Across approximately 2,000 logged driver days, drivers tend to stop work earlier when they have been having relatively good days in terms of fare volume. This results in conspicuously suboptimal earnings over the course of a year because taxi drivers are deciding to knock off early on exactly the days that are the most lucrative. By rational decision-making accounts, the days with high wages per hour are exactly the days that drivers should be exploiting by working for a long time. Although suboptimal, this behavior is consistent with decision making via an “aspiration level” heuristic via which a driver stops working in a day once a certain threshold amount of money has been earned. These fare logs suggest that taxi drivers are thinking of their earnings on a short, day-by-day basis rather than setting a longer term policy that would improve their multi-day wages per hour ratio.

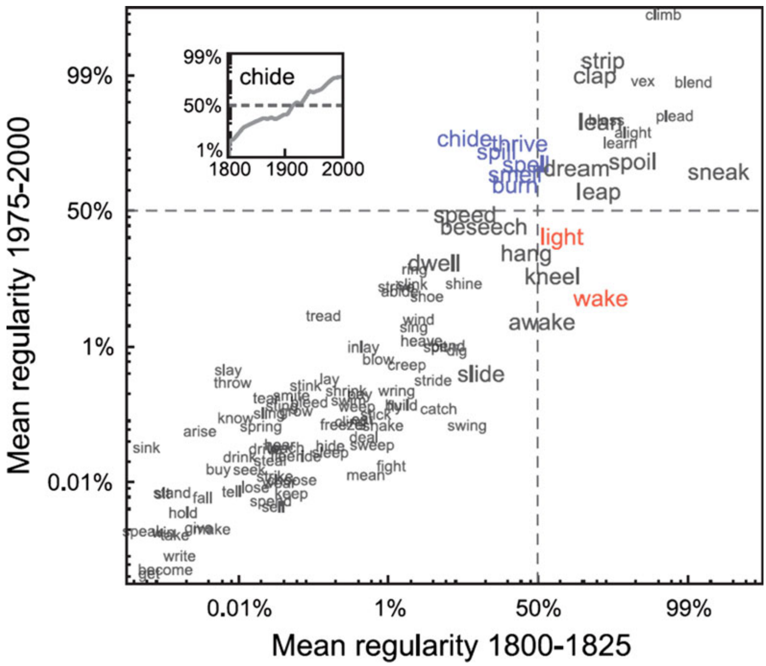

Fig. 2. A plot from Michel et al. (2011) of the mean regularity of the verbs in the period 1800–1825 versus their mean regularity from 1975 to 2000. For example, the irregular form "fought" has been preferred over "fighted" as the past tense of "fight" for the both time periods, with the regular form "fighted" being used about 1% of the time 1800–1825 and about 0.02% of the time 1975–2000. The inset shows the gradual increase in the regularity of the verb "chide" over 200 years. In the 1800s, "chid" was far more common than "chided," but exactly the opposite is true nowadays. The verbs shown in blue have switched from being predominantly irregular in 1800–1825 to being predominantly regular in 1975–2000, whereas the verbs shown in red have made the opposite transition. Figure reprinted with permission from the authors.

4.4. Zhang (2010)

A life-changing decision that individuals with poor kidney functioning must make is whether to accept a kidney that is being offered to them from a donor, typically with a hospital acting as intermediary. There is a waiting list of candidate recipients for kidneys that can be quite long. A kidney, when it become available, is offered to the first candidate on the waiting list, and if they decline it (often due to idiosyncratic timing and health considerations), it is then offered to the second candidate, and so on. The decisions of candidate recipients show evidence of strong observational learning effects. In particular, candidates later on the waiting list use the earlier candidates' decisions to decline a kidney as evidence that the kidney may be generally unsuitable. Fig. 4 shows a striking example of this effect looking at a subset of the data for which two kidneys can be equated for their clinical utility because they are donated from the same, deceased individual. Zhang divided these kidneys into two matched groups—those that were accepted earlier from their matched pair and those that were accepted later. Fig. 4 shows strong deviations in

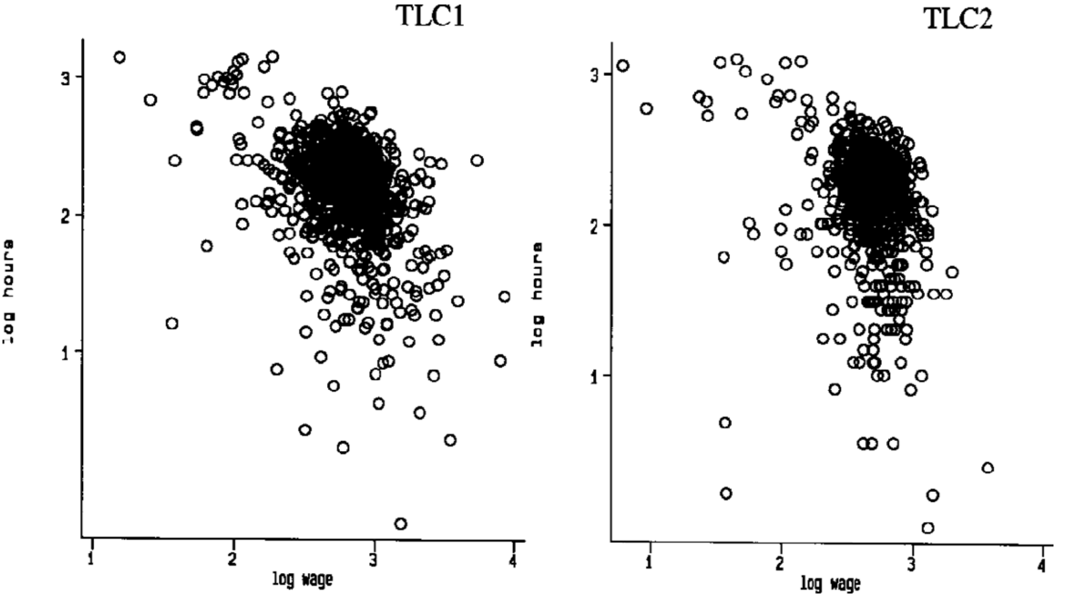

Fig. 3. Camerer et al.'s (1997) study of fare logs from New York City taxi drivers. The results from two separate data sets, TLC1 and TLC2, show that as hourly wages increase, the total number of hours worked on that day decrease. Figure reprinted with permission from the authors.

the profile of acceptance for the matched kidneys, indicated by the wide gaps between adjacent bars for matched kidneys that came from the same individual. Once a kidney has been declined by one candidate recipient, it becomes more likely to be declined by subsequent candidates—candidates who are aware of the history of rejections of a particular kidney that they are being offered. Statistical modeling indicates that the number of previous rejections of a kidney has a strong negative relation on the likelihood that a kidney will be accepted by a candidate recipient. Despite a continual shortage of kidneys, a self-reinforcing chain of inferences about the quality of kidney is formed (e.g., "I'm suspicious of a kidney that has been declined by 20 people ahead of me. Thus, I will also decline it."). There is a robust literature on the unique prevalence and power of observational learning in humans (Dean, Kendal, Shapiro, Thierry, & Laland, 2012; Goldstone *et al.*, 2013; Wisdom, Song & Goldstone, 2013), but this example from the U.S. kidney market is particularly powerful because it documents that people imitate others' preferences for an object even when they are literally dying for the object.

4.5. Khatib et al. (2011)

Game play offers a fertile research paradigm for cognitive science, as evidenced by a forthcoming *topiCS* topic dedicated to action games in cognitive science (Gray, 2016). One area of cognitive science interest in game play is "serious games"—games that have a purpose other than entertainment alone, particularly education, training, or solving real-

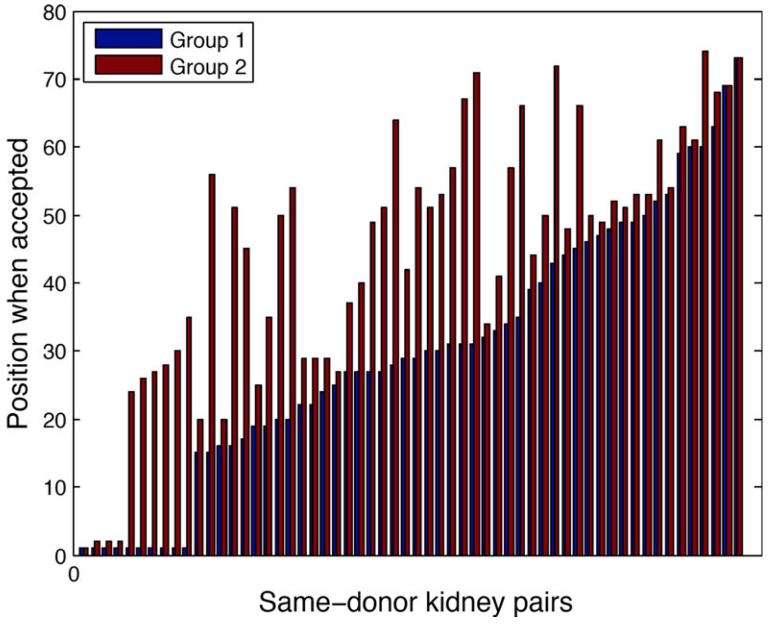

Fig. 4. Results from Zhang (2010) on acceptances of kidneys by candidates for kidney transplants. This graph shows the subset of results when two kidneys of equivalent clinical utility are donated by the same, deceased individual. Group 1 contains 58 kidneys that were accepted earlier in the queue of candidates. Group 2 shows the matched kidneys that were accepted later in the queue. Kidneys are matched when they come from the same individual donor, and they are plotted adjacent to each other. The vertical axis shows the position on the queue of the kidney when it was accepted. For example, for the fifth pair of kidneys from the left, one kidney was accepted by the first candidate it was offered to, but the other kidney was not accepted until it had been declined by 22 candidates. Many matched kidneys differ substantially in terms of when they are accepted, consistent with the notion that candidate kidney recipients use the rejection of a kidney by candidates ahead of them in queue as evidence of the kidney's low quality. Figure reprinted with permission from the authors.

world problems (Mayer, 2014). One notable example of a serious game is Foldit (https://fold.it/portal/), which tasks its participants with folding proteins into as low an energy state as possible, a notoriously difficult problem even for the most sophisticated artificial intelligence systems available (Cooper et al., 2010). Scientists can analyze the best solutions found by players to determine if they can be applied to understanding or manipulating proteins in the real world. For example, in 2011, Foldit players uncovered the crystal structure of a virus that causes AIDS in monkeys, producing a solution that had eluded professional scientists for 15 years. In 2012, using a version of the game that allows for the creation of new proteins, game players constructed an enzyme that can speed up a biosynthetic reaction used in a variety of drugs, including cholesterol medications, by 2,000%. Khatib et al. (2011) studied the strategies that 57,000 Foldit players use to achieve these successes. A key to players' success is the strategic use of algorithmic "recipes" to supplement their human search. Players, in effect, turn themselves into cyborg search agents, combining human intuition and the computational efficiency of crafted and tailored algorithms. Fig. 5 shows the prevalence of deploying different classes of algorithms at different points in a solution by different players. The third category of recipes, "local optimize," performs local energy minimizations along the protein backbone. All players use this strategy increasingly often as their work on a protein progresses, and that trend is particularly pronounced for the best players. The top players seem to appreciate that applying this strategy on the initial state of a Foldit puzzle is not effective because a successful prediction for how a protein will fold will no longer have the general backbone shape of the initial state. This kind of insight can be harvested for "algorithm mining" efforts to improve machine learning by incorporating successful human strategies, and it also speaks more generally to the global-to-local search processes that successful problem solvers employ.

5. What NODS are and are not about

In assessing the opportunities for NODS, it is helpful to clarify what we believe the articles in this issue to be about, and not about. Beginning with what this topic is *not* about is arguing for a "bigger is better" ideology. In particular, we have eschewed framing this topic in terms of "Implications of Big Data for Cognitive Science" despite the current zeitgeist surrounding "Big Data." Bigger is not necessarily better when it comes to data (e.g., Roberts & Winters, 2013 for discussion). Many computer scientists interested in Big Data are interested in developing technologies that allow users to process tera-, peta-, and exabytes of data. However, some of the data sets that have been most psychologically revealing, like John Anderson's emails and taxi drivers' logs, are mere megabytes or less. The work of Camerer et al. on taxi drivers' work hours is important for cognitive science not because it requires high-speed computers or new data-mining algorithms—it does not. It is important for what it reveals about real drivers' decisions to

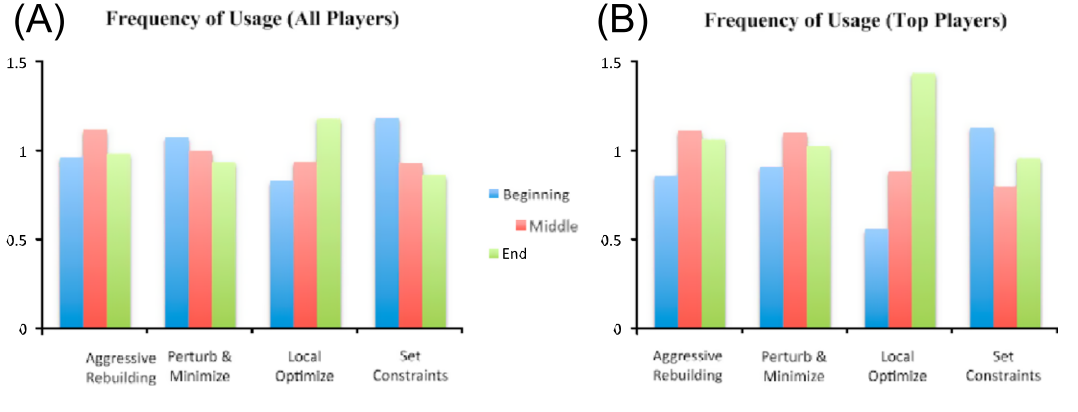

Fig. 5. Use of strategies by Foldit players as documented by Khatib et al. (2011). Panel A shows the frequency of different strategies for all players, while Panel B shows the same frequencies for the top performing players. The "local optimize" strategy is used increasingly often as the problem solving effort extends over time, and this trend is particularly pronounced for the best players. This strategy performs local energy minimizations along the protein backbone at every segment of the protein. Figure reprinted with permission from the authors.

stop driving for the day, and why a driver's manifest use of an "aspiration level" decision rule personally costs him/her thousands of dollars of wages. Although it runs against the grain of some sectors of the Big Data movement, we nonetheless maintain that a quality piece of data is often worth 1,000 times its weight in raw and unwashed data.

This topic is also not concerned with using large data sets to reveal community, institutional, or social network patterns. This work is fascinating and important, and we firmly believe in the relevance of group-level patterns and processes for cognitive science (Goldstone & Gureckis, 2009). However, the articles in this topic are focused more on revealing principles of individual psychology through NODS. Sociology and social network analysis are large, additional topics worthy of their own coverage. Plus, it is perhaps more eye-opening for cognitive scientists to think about ways in which their own, typically individual-oriented domains of inquiry into human cognition can be transformed by examining natural data sets.

Another enterprise this topic is not about is developing and using tools for visualizing, storing, or analyzing large amounts of experimentally collected data, such as fMRI, EEG, body movements, and video data. These high-bandwidth data sources fall within the purview of Big Data for Cognitive Science, but not NODS because they are collected in the context of laboratory experiments. A final undertaking this topic is not about is trying to convince psychologists not to run experiments. All of the articles in this topic provide examples of the rich possibilities for interplay between experiments and NODS.

Moving on to the positive characterization of our interpretation of NODS for cognitive science, for us the real excitement is not necessarily the size of the data sets, but their availability. The increasing openness of data makes it possible for people to look for patterns in data sets that never would have occurred to the individuals who originally created the data sets. Inspecting naturally occurring data sets is very much in the spirit of native North American Inuits alleged use of the whole whale—meat for food, blubber for candle oil, bones for sleds, and skin for clothing. Data sets that are made available either for free or for a price have often been developed at great expense, and hence have been explicitly built by their creators so that a variety of people can plumb the data for their own purposes. There are many evocative, revealing data sets that can help inform our theories of cognition. Empowered with currently available tools in database querying, automatic structuring of data, statistical analysis, and visualization, the worn cliché is apt: we really are only limited by our creativity. The articles in this topic demonstrate that theoretical questions should be driving the research agenda. Questions about how intelligent systems in rich environments make decisions, remember information, perceive objects, learn language, and adapt to solve new problems are the primary focus. However, given this agenda, NODS are serving as theory accelerators. The theories we are able to advance are much more nuanced than possible without the NODS, and they have dramatically accelerated the cycle of theory development, theory testing, and theory revision.

The five previous case studies and the current articles provide inspirational examples of researchers who have creatively interrogated existing data sets in surprising and genuinely novel ways. Cognitive scientists, equipped with powerful theories of cognition, are well positioned to add new perspectives on what to look for in existing data sets. For example, to a chess player, the massive repository of online chess games is almost exclusively viewed from the restricted vantage point of an historic record of games that can possibly be scrutinized for strategic insights. However, to a cognitive scientist steeped in the problem solving literature, they represent a treasure trove of information about how factors like blindfolding, experience, and speed affect the quality of thinking, and the roles of search versus perceptual matching for problem solving (Chabris & Hearst, 2003). As Karmiloff-Smith and Inhelder (1975) have argued, "To get ahead, get a theory." This is just as true of modern cognitive scientists as it is of the children that Karmiloff-Smith studied. The rapidly increasing availability of high-quality data logging, storage, querying, visualization, and analysis tools means that the bottleneck for progress in the social sciences is increasingly on the theoretical side. Cognitive science has always been strong on theory, and its practitioners have unique theoretical vantage points from which to explore real-world data sets. If theory is indeed the bottleneck, then cognitive science becomes a perspective of increasing societal relevance because of its ability to widen the constriction.

Acknowledgments

We thank Colin Camerer, Rick Dale, Simon Dedeo, Wayne Gray, Michael Jones, Firas Khatib, David Landy, Jean-Baptiste Michel, Lael Schooler, Peter Todd, Amanda Windburn, Juanjuan Zhang, and the authors of the articles in this topic for helpful comments, figures, conversations, and advice.

References

Adamic, L. A., & Glance, N. (2005). The political blogosphere and the 2004 US election: Divided they blog. *Proceedings of the 3rd international workshop on Link discovery*, pp. 36–43. New York.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, *2*, 396–408.

Anderson, J. R., & Schooler, L. J. (2000). The adaptive nature of memory. In E. Tulving & F. I. M. Craik (Eds.), *Handbook of memory* (pp. 557–570). New York: Oxford University Press.

Beilock, S. L., & Gray, R. (2007). Why do athletes "choke" under pressure? In G. Tenenbaum & R. C. Eklund (Eds.), *Handbook of sport psychology* (3rd ed., pp. 425–444). Hoboken, NJ: John Wiley & Sons.

Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, *25*, 1798–1828.

Berger, J., Bradlow, E. T., Braunstein, A., & Zhang, Y. (2012). From Karen to Katie: Using baby names to study cultural evolution. *Psychological Science*, *23*, 1067–1073.

Berger, J., & Le Mens, G. (2009). How adoption speed affects the abandonment of cultural tastes. *Proceedings of the National Academy of Sciences*, *106*, 8146–8150.

Berger, J., & Pope, D. (2011). Can losing lead to winning? *Management Science*, *57*, 817–827.

Bernartzi, S., & Thaler, R. H. (2007). Heuristics and biases in retirement savings behavior. *Journal of Economic Perspectives*, *21*, 81–104.

Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D. I., Marlow, C., Settle, J. E., & Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, *489*, 295–298.

Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, *105*(38), 14325–14329.

Braithwaite, D. W., Goldstone, R. L., van der Maas, H. L. J., & Landy, D. H. (2016). Informal mechanisms in mathematical cognitive development: The case of arithmetic. *Cognition*, *149*, 40–55.

Camerer, C., Babcock, L., Loewenstein, G., & Thaler, R. (1997). Labor supply of New York City cabdrivers: One day at a time. *Quarterly Journal of Economics*, *112*, 407–441.

Castranova, E. (2005). *Synthetic worlds*. Chicago: University of Chicago Press.

Chabris, C. F., & Hearst, E. S. (2003). Visualization, pattern recognition, and forward search: Effects of playing speed and sight of the position on grandmaster chess errors. *Cognitive Science*, *27*, 637–648.

Clark, A. (2004). *Natural born cyborgs*. Oxford, England: Oxford University Press.

Clark, A. (2009). *Supersizing the mind*. Oxford, England: Blackwell.

Cooke, N. J., & Hilton, M. L. (Eds.) (2015). *Enhancing the effectiveness of team science*. Washington, DC: National Academies Press.

Cooper, S., Khatibi, F., Treuille, A., Barbero, J., Lee, J., Beenen, M., Leaver-Fay, A., Baker, D., Popovic, Z., & Foldit Players (2010). Predicting protein structures with a multiplayer online game. *Nature*, *466*, 756–760.

Dean, L. G., Kendal, R. L., Shapiro, S. J., Thierry, B., & Laland, K. N. (2012). Identification of the social and cognitive processes underlying human cumulative culture. *Science*, *335*, 1114–1118.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, *47*(2), 315–372.

Ebbinghaus, H. (1885). *Memory: A contribution to experimental psychology*. New York: Dover.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.

Goldstone, R. L., & Gureckis, T. M. (2009). Collective behavior. *Topics in Cognitive Science*, *1*, 412–438.

Goldstone, R. L., Wisdom, T. N., Roberts, M. E., & Frey, S. (2013). Learning along with others. *Psychology of Learning and Motivation*, *58*, 1–45.

Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, *37*, 424–438.

Gray, W. (2016). Action games as a cognitive science paradigm. Topics in Cognitive Science.

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. T. (2007). Topics in semantic representation. *Psychological Review*, *114*(2), 211–244.

Gureckis, R. L., & Goldstone, R. L. (2009). How you named your child: Understanding the relationship between individual decision making and collective outcome. *Topics in Cognitive Science*, *1*, 651–674.

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

Jones, M. N. (Ed.) (2016). *Big data in cognitive science: From methods to insights*. New York: Psychology Press.

Karmiloff-Smith, A., & Inhelder, B. (1975). If you want to get ahead, get a theory. *Cognition*, *3*, 195–212.

Khatibi, F., Cooper, S., Tyka, M. D., Xu, K., Macedon, I., Popovic, Z., Baker, D., & Foldit Players (2011). Algorithm discovery by protein folding game players. *Proceedings of the National Academy of Sciences*, *108*, 18949–18953.

Koedinger, K. R., Stamper, J. C., Leber, B., & Skogsholm, A. (2013). LearnLab's DataShop: A data repository and analytics tool set for cognitive science. *Topics in Cognitive Science*, *5*(3), 668–669.

Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, *111*, 8788–8790.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*, 436–444.

Lupyan, G., & McClelland, J. L. (2003). Did, made, had, said: Capturing quasi-regularity in exceptions. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 740–745). Mahwah, NJ: Lawrence Erlbaum.

Mace, W. M. (1977). James J. Gibson's strategy for perceiving: Ask not what's inside your head, but what your head's inside of. In R. E. Shaw & J. Bransford (Eds.), *Perceiving, acting, and knowing* (pp. 43–63). Hillsdale, NJ: Erlbaum.

Madrian, B., & Shea, D. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics*, *116*(4), 1149–1187.

Mayer, R. E. (2014). *Computer games for learning: An evidence-based approach*. Cambridge, MA: MIT Press.

Michel, J., Shen, Y. K., Aiden, A. P., Veres, A., Gray, M. K., Team, T. G. B., Pickett, J. P., Holberg, D., Clancy, D., Norvig, P., Orwant, J., Pinker, S., Nowak, M. A., & Aiden, E. L. (2011). Quantitative analysis of culture using millions of digitized books. *Science*, *331*, 176–182.

Monaghan, P., Christiansen, M. H., & Fitneva, S. A. (2011). The arbitrariness of the sign: Learning advantages from the structure of the vocabulary. *Journal of Experimental Psychology: General*, *140*, 325–347.

Nathan, M. J., & Alibali, M. W. (2010). Learning sciences. *Wiley Interdisciplinary Reviews—Cognitive Science*, *1*, 329–345.

Nathan, M. J., & Sawyer, K. (2014). Foundations of Learning Sciences. In K. Sawyer (Ed.), *The Cambridge handbook of the learning sciences* (2nd ed., pp. 21–43). Cambridge, England: Cambridge University Press.

Neisser, U. (1976). *Cognition and reality: Principles and implications of cognitive psychology*. New York: Freeman.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Mateo, CA: Morgan Kaufmann.

Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge, England: Cambridge University Press.

Pearl, J. (2009). Causal inference in statistics: An overview. *Statistics Surveys*, *3*, 96–146.

Pope, D. G., & Schweitzer, M. E. (2011). Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes. *American Economic Review*, *101*, 129–157.

Post, T., Assem, M. J., Baltussen, G., & Thaler, R. H. (2008). Deal or no deal? Decision making under risk in a large-payoff game show. *American Economic Review*, *98*, 38–71.

Roberts, S., & Winters, J. (2013). Linguistic diversity and traffic accidents: Lessons from statistical studies of cultural traits. *PLoS ONE*, *8*(e70902), 1–13.

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, *311*, 854–856.

Shepard, R. N. (1984). Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming. *Psychological Review*, *91*, 417–447.

Steyvers, M., & Tenenbaum, J. B. (2005). Graph theoretic analyses of semantic networks: Small worlds in semantic networks. *Cognitive Science*, *29*, 41–78.

Sugihara, G., May, R., Ye, H., Hsieh, C., Deyle, E., Fogerty, M., & Munch, S. (2012). Detecting causality in complex ecosystems. *Science*, *338*, 496–500.

Szell, M., Lambiotte, R., & Thurner, S. (2010). Multirelational organization of large-scale social networks in an online world. *Proceedings of the National Academy of Sciences*, *107*, 13636–13641.

Wisdom, T. N., Song, X., & Goldstone, R. L. (2013). Social learning strategies in a networked group. *Cognitive Science*, *37*, 1383–1425.

Zhang, J. (2010). The sound of silence: Observational learning in the U.S. kidney market. *Marketing Science*, *29*, 315–335.

Articles in this Topic

Berger, J. (2016). Does presentation order impact choice after delay? *Topics in Cognitive Science*.

Christiansen, M. H., & Monaghan, P. (2016). Division of labor in vocabulary structure: Insights from corpus analyses. *Topics in Cognitive Science*, *8*, 670–684. doi: 10.1111/tops.12205.

Griffiths, T. L., Abbott, J. T., & Hsu, A. S. (2016). Exploring human cognition using large image databases. *Topics in Cognitive Science*, *8*, 569–588. doi: 10.1111/tops.12209.

Heit, E., & Nicholson, S. P. (2016). Missing the party: Political categorization and reasoning in the absence of party label cues. *Topics in Cognitive Science*, *8*, 697–714. doi: 10.1111/tops.12206.

Koedinger, K. R., Yudelson, M. V., & Pavlik Jr., P. I. (2016). Testing theories of transfer using error rate learning curves. *Topics in Cognitive Science*, *8*, 589–609. doi: 10.1111/tops.12208.

Moat, H. S., Olivola, C. Y., Chater, N., & Preis, T. (2016). Searching choices: Quantifying decision making processes using search engine data. *Topics in Cognitive Science*, *8*, 685–696. doi: 10.1111/tops.12207.

Pope, D. G. (2016). Exploring psychology in the field: Steps and examples from the used-car market. *Topics in Cognitive Science*, *8*, 660–669. doi: 10.1111/tops.12210.

Vincent-Lamarre, P., Blondin Massé, A., Lopes, M., Lord, M., Marcotte, O., & Harnad, S. (2016). The latent structure of dictionaries. *Topics in Cognitive Science*, *8*, 625–659. doi: 10.1111/tops.12211.