



Original Articles

Exploration and exploitation of Victorian science in Darwin's reading notebooks

Jaimie Murdock^{a,b}, Colin Allen^{a,c,d}, Simon DeDeo^{a,b,e,f,*}^a Program in Cognitive Science, Indiana University, Bloomington, IN 47405, USA^b School of Informatics and Computing, Indiana University, 919 E. 10th Street, Bloomington, IN 47408, USA^c Department of History and Philosophy of Science and Medicine, Indiana University, Bloomington, IN 47405, USA^d School of Humanities and Social Sciences, Xi'an Jiaotong University, Xi'an, China^e Department of Social and Decision Sciences, Carnegie Mellon University, 5000 Forbes Avenue, BP 208, Pittsburgh, PA 15213, USA^f Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

ARTICLE INFO

Article history:

Received 10 March 2016

Revised 23 November 2016

Accepted 28 November 2016

Keywords:

Cognitive search

Information foraging

Topic modeling

Exploration-exploitation

History of science

Scientific discovery

ABSTRACT

Search in an environment with an uncertain distribution of resources involves a trade-off between exploitation of past discoveries and further exploration. This extends to information foraging, where a knowledge-seeker shifts between reading in depth and studying new domains. To study this decision-making process, we examine the reading choices made by one of the most celebrated scientists of the modern era: Charles Darwin. From the full-text of books listed in his chronologically-organized reading journals, we generate topic models to quantify his local (text-to-text) and global (text-to-past) reading decisions using Kullback-Liebler Divergence, a cognitively-validated, information-theoretic measure of relative surprise. Rather than a pattern of surprise-minimization, corresponding to a pure exploitation strategy, Darwin's behavior shifts from early exploitation to later exploration, seeking unusually high levels of cognitive surprise relative to previous eras. These shifts, detected by an unsupervised Bayesian model, correlate with major intellectual epochs of his career as identified both by qualitative scholarship and Darwin's own self-commentary. Our methods allow us to compare his consumption of texts with their publication order. We find Darwin's consumption more exploratory than the culture's production, suggesting that underneath gradual societal changes are the explorations of individual synthesis and discovery. Our quantitative methods advance the study of cognitive search through a framework for testing interactions between individual and collective behavior and between short- and long-term consumption choices. This novel application of topic modeling to characterize individual reading complements widespread studies of collective scientific behavior.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

The general problem of “information foraging” (Pirulli & Card, 1999) in an environment about which agents have incomplete information has been explored in many fields, including cognitive psychology (Hills, Todd, Lazer, Redish, & Couzin, 2015; Todd et al., 2012), neuroscience (Cohen, McClure, & Yu, 2007), economics (Azoulay-Schwartz, Kraus, & Wilkenfeld, 2004; March, 1991), finance (Uotila, Maula, Keil, & Zahra, 2009), ecology (Eliassen, Jørgensen, Mangel, & Giske, 2007; Stephens & Krebs, 1986), and computer science (Sutton & Barto, 1998). In all of these areas, the

searcher aims to enhance future performance by surveying enough of existing knowledge to orient themselves in the information space.

Individual scientists and scholars can be viewed as conducting a cognitive search (Todd et al., 2012) in which they must balance *exploration* of ideas that are novel to them against *exploitation* of knowledge in domains in which they are already expert (Berger-Tal, Nathan, Meron, & Saltz, 2014). Researchers have studied the exploration-exploitation trade-off in cognitive search at timescales of minutes up to years and decades. Laboratory experiments on visual attention are one example of this balancing act on short timescales (Chun & Wolfe, 1996), while studies of the recombination of patented technologies demonstrate long-term group behavior (Youn, Strumsky, Bettencourt, & Lobo, 2015). New advances in the digitization of historical archives enable longitudinal study of

* Corresponding author at: Department of Social and Decision Sciences, Carnegie Mellon University, 5000 Forbes Avenue, BP 208, Pittsburgh, PA 15213, USA.

E-mail addresses: jammurdo@indiana.edu (J. Murdock), colallen@indiana.edu (C. Allen), simon@santafe.edu (S. DeDeo).

how individuals explore and synthesize the work of their contemporaries and predecessors over the course of a lifetime.

As one of the most successful and celebrated scientists of the modern era, Charles Darwin's scientific creativity has been the subject of numerous narrative and qualitative studies (Gruber & Barrett, 1974; Johnson, 2010; Van Hulle, 2014). In part, these studies are possible because Darwin left his biographers careful records of his intellectual and personal life. These include records of the books he read from 1837 to 1860, a critical period which culminated in the publication of *The Origin of Species*. Table 1 summarizes key events in Darwin's life.

This article presents the first quantitative analysis of an important scientist's reading diaries, tracking how Darwin navigated the exploration-exploitation trade-off in choosing what to read. We link Darwin's reading records with the full text of the original volumes. We then use probabilistic topic models (Blei, Ng, & Jordan, 2003; Blei, 2012b) to represent the original text of each book Darwin read as a mixture of topics. We use information theory to measure the surprise, or unpredictability, of the next book that Darwin chose to read, compared to his past history of reading.

We present three key findings:

1. Darwin's reading patterns switch between both exploitation and exploration throughout his career. This is in contrast to a pure surprise-minimization strategy that consistently exploits content within a local region before moving on. The general trend, as Darwin's career develops, is towards increasing exploration.
2. In comparison to the publication order of the texts Darwin read, Darwin's reading order shows higher average surprise. This indicates that the order in which the books were written by the scientific community is less surprising than the order in which Darwin read them.
3. Darwin's strategies fall into three long-term epochs, or behavioral modes characterized by distinct patterns of surprise-seeking. These epochs correspond to three biographically significant periods: Darwin's post-*Beagle* studies, his extensive work on barnacles, and a final period leading to his synthesis of natural selection in the *Origin of Species*.

While the bulk of empirical studies in cognitive science are concerned with measuring population-level effects due to experimental manipulations, case studies play an important role in driving cognitive theorizing, experimentation, and modeling. For example, the case of the memory-impaired patient H.M. has driven many advances in cognitive neuroscience and computational models of memory (reviewed in Squire & Wixted (2011)). Other case studies, such as that of the frontal-lobe injury in Phineas Gage, provide important contrasts for later studies (reviewed in Macmillan (2000)).

Detailed longitudinal investigations of a single individual may involve repeated trials and changing strategies that cannot be observed in laboratory experiments involving a single task. Cognitive science could be enriched by using longitudinal studies such as ours to design laboratory studies with higher ecological validity. However, it is a challenging task to design laboratory studies of (for example) reading choices that take into account a subject's extensive prior history of reading decisions.

An advantage of taking Charles Darwin as a case study is the extensive attention that he has received from historians and biographers since his death. These studies provide a novel means for validating our mathematical tools. We will present a number of theoretical and empirical reasons why our methods measure cognitively-relevant features of a reader's experience. At the same time, the fact that our results also recover key features of Darwin's

Table 1

Timeline. Major events in Charles Darwin's life, including those marked in Fig. 1. This paper focuses on the critical period of his work from 1837 to 1860, leading to the publication of *The Origin of Species* (boundaries marked in bold). See Berra (2009) for an expanded chronology.

Major Events in Charles Darwin's Life (1809–1882)	
12 February 1809	Born in Shrewsbury, England
22 October 1825	Matriculates at University of Edinburgh
15 October 1827	Admitted to Christ's College, Cambridge
27 December 1831	Departs England aboard the <i>HMS Beagle</i>
2 October 1836	Return to England aboard the <i>HMS Beagle</i>
July 1837	First entries in reading notebooks
August 1839	Publication of <i>The Voyage of the Beagle</i> (1st edition)
May 1842	Writes the 1st Essay on Species
4 July 1844	Writes the 2nd Essay on Species
August 1845	Publication of <i>The Voyage of the Beagle</i> (2nd edition)
1 October 1846	Begins barnacle project
19 February 1851	Publishes first volume of barnacle work
9 September 1854	Begins sorting notes on natural selection
14 May 1856	Starts writing "large work" on species
24 November 1859	Publication of <i>The Origin of Species</i> (1st edition)
13 May 1860	Last entry in reading notebooks
24 February 1871	Publication of <i>The Descent of Man</i>
19 February 1872	Publication of <i>The Origin of Species</i> (6th and final edition)
21 April 1882	Dies at Down House in Kent, England

intellectual life provides additional support for the reliability of our methods.

We cannot claim Darwin's information foraging behavior is typical for either scientists or the population-at-large; however, our analytic methods may be applied to other case studies to find population-level generalizations, provided there is access to the full text of each reading.

Our approach contrasts with previous uses of topic modeling to analyze the large-scale structure of scientific disciplines (Blei & Lafferty, 2007; Cohen Priva & Austerweil, 2015; Griffiths & Steyvers, 2004; Hall, Jurafsky, & Manning, 2008) and the humanities (Blei, 2012a; Jockers, 2013; Mohr & Bogdanov, 2013), which are each created through the collective effects of individual-level behavior. Previous models of historical records have focused on language use as an indication of larger shifts in style (Hughes, Foti, Krakauer, & Rockmore, 2012; Underwood & Sellers, 2012), learnability (Hills & Adelman, 2015), or content (Goldstone & Underwood, 2014; Klingenstein, Hitchcock, & DeDeo, 2014; Michel et al., 2011) of significant portions of publications in a field, including a study of *Cognition* itself (Cohen Priva & Austerweil, 2015).

These works model the collective state of all published works at a particular date, but obscure the role of individual foraging behavior. By focusing on a single individual for whom ample records exist, we gain access to what Tria, Loreto, Servedio, and Strogatz (2014) describe as "the interplay between individual and collective phenomena where innovation takes place".

2. Materials and methods

2.1. Darwin's reading notebooks

Darwin was a meticulous record-keeper—starting in April 1838, he kept a notebook of "books to be read" and "books read". These records span the 23 years from 1837 to 1860, tracking his reading choices from just after his return to England aboard the *HMS Beagle* to just after the publication of *The Origin of Species*. We located the full-text of 665 of the 687 (96.7%) English non-fiction mentioned in these reading notebooks through a variety of online digital libraries. See Appendix A for details of corpus curation.

2.2. Probabilistic topic models

We model these texts using Latent Dirichlet Allocation (LDA; Blei et al. (2003), Blei (2012b)), a type of probabilistic topic model. LDA is a generative model that represents each document as a bag of words generated by a mixture of topics. LDA topic models have been extensively used to analyze human-generated text across many domains (Blei, 2012a; Blei & Lafferty, 2007; Griffiths & Steyvers, 2004; Jockers, 2013; Mohr & Bogdanov, 2013). Furthermore, topic models predict the behavior of human subjects in a variety of word association and disambiguation tasks (Griffiths, Steyvers, & Tenenbaum, 2007). Topic models allow us to describe Darwin's reading as taking place in a $(k - 1)$ -dimensional space, the simplex, where a particular volume is described as a probability distribution, \vec{p} , over k topics. To test the robustness of our results, we vary the number of topics, k . In the main body of the paper, we report for $k = 80$ because we found that our results were robust to alternative values of k and to differing random seeds for each; see Appendix F for $k = \{20, 40, 60\}$. Darwin's "semantic voyage" is the path he takes through this space from text to text.

2.3. Cognitive surprise and Kullback-Leibler divergence

We use probabilistic topic models to quantify the structure of the texts. Characterizing Darwin's reading choices then becomes the problem of comparing distributions over topics. Our goal is to quantify surprise: the extent to which a new reading either satisfies, or violates, the expectations a reader would have based on what came before. A high-surprise reading indicates that the reader has chosen a book that contains a new mixture of topics compared to what came before.

In our model, \vec{p} and \vec{q} are topic distributions over texts or an aggregation of texts. The reader is modeled in an information-theoretic fashion as building efficient mental representations of these distributions. Our task is then to quantify his exploratory behavior as he moves from one distribution to another.

To compare two distributions in this way we use the Kullback-Leibler (KL) divergence, first introduced by Kullback and Leibler (1951). It is defined by

$$D_{\text{KL}}(\vec{q} \parallel \vec{p}) = \sum_{i=1}^k q_i \log_2 \frac{q_i}{p_i}, \quad (1)$$

where \vec{p} is the distribution over topics that the reader encountered before, and \vec{q} the new distribution of topics that the reader encounters next. (In this paper, we consider two choices for \vec{p} : (1) the distribution over topics for the just previous book, and (2) the average over all books in the reader's past. Meanwhile, \vec{q} will always refer to a distribution over topics for the next book a reader encounters.)

KL divergence can be thought of as a measure of surprise associated with a change in representation of a learner who encounters the unexpected in ways that require her to alter the way she represents the world. An agent expecting observations to be drawn from probability distribution \vec{p} will have those expectations violated if it arrives according to a different distribution \vec{q} . KL divergence quantifies the surprise in a particularly natural fashion; to aid the reader unfamiliar with this quantity, we present two ways to understand the mathematical structure of KL in Appendix B: first in terms of "coding failure" and second in terms of the rate of accumulation of Bayesian evidence against one distribution and in favor of the other.

The use of KL to measure cognitive surprise has had a number of recent successes. Vision researchers have used KL to track surprise

in a visual scene: the places on the screen where new events most violate the viewer's previous assumptions; KL then accurately captures attention attractors in visual search tasks (Demberg & Keller, 2008; Itti & Baldi, 2009). More generally, KL has seen wide use in the cognitive sciences; Resnik (1993), for example, proposed the KL divergence as a measure of selectional preferences in language (reviewed in Light & Greiff (2002)). It has found use in many successful models of linguistic discrimination, including syntactic comprehension (Hale, 2001; Levy, 2008), speech recognition (Calamaro & Jarosz, 2015; Martin, Peperkamp, & Dupoux, 2013) and word sense disambiguation (Resnik, 1997).

2.4. Text-to-text and text-to-past surprise

We use KL in two distinct ways. We measure the text-to-text surprise: given a distribution over topics for the text Darwin just read, how surprised is he upon encountering the next volume's topic distribution? Text-to-text surprise is a *local* measure. We also measure the text-to-past surprise: given all of the volumes that Darwin has encountered so far, how surprised is Darwin by the text that comes next? Text-to-past surprise is a *global* measure.

If we define θ_i as the topic distribution for document i , then text-to-text (T2T) and text-to-past (T2P) surprise are defined as

$$T2T(i) = D_{\text{KL}}(\theta_i \parallel \theta_{i-1}), \quad (2)$$

$$T2P(i) = D_{\text{KL}}\left(\theta_i \parallel \frac{\sum_{j=0}^{i-1} \theta_j}{i}\right). \quad (3)$$

Text-to-text surprise and text-to-past surprise provide complementary windows onto Darwin's decision-making. Local decision-making, meaning the choice of the next text to read given the current one, is captured by text-to-text surprise. Global decision-making, the choice of which text to read given the entire history of reading to date, is captured by text-to-past surprise. Low surprise, in either case, is a signal of *exploitation*, while high surprise indicates larger jumps to lesser-known topics, and thus of *exploration*. These measures can be easily generalized to arbitrary text-to- N surprise measures, representing the choice of the next reading given the history of readings within the past N volumes or time periods.

We characterize Darwin's decision process by the combination of text-to-text and text-to-past surprise. Exploration, indicated by high surprise, happens when a searcher is moving across a space not previously explored. Exploitation, indicated by low surprise, happens when a searcher has a sustained focus on material they are already familiar with.

These local and global behaviors do not have to align. For example, text-to-text surprise may be high (local exploration) at the same time that text-to-past surprise is low (global exploitation). This can happen if Darwin's readings interleave different topics that he has already seen. If, for example, Darwin alternates between readings in philosophy with readings in travel narratives, then each local jump will have high KL (a travel narrative is dominated by topics that are rare in a philosophical text, and vice versa) and Darwin's readers will appear as a local exploration. However, once this pattern of alternation has been established, the average over past texts will include both philosophical and travel narrative topics, lowering the text-to-past surprise, and driving the system back towards global exploitation.

Conversely, text-to-text surprise can also be low while text-to-past surprise is high. This can happen if Darwin has recently begun a novel, but focused, investigation. In this situation, he focuses on a particular subset of topics that are under-represented in his overall history. If Darwin begins by focusing on philosophical texts, and then switches to travel narratives, his second (and subsequent)

travel narrative readings will have low text-to-text surprise; but the average over past texts will be dominated by a long history of philosophical readings, leaving the text-to-past surprise high until he has accumulated so many readings of the latter type that they dominate the past average.

2.5. Cultural production and null reading models

Our methods also allow us to understand how Darwin selectively re-orders the products of his culture. To do this, we measure text-to-text and text-to-past surprise using the publication order of the texts that Darwin read. We can then ask questions about the extent to which Darwin's encounters with the products of his culture were more or less exploratory than the order in which they were produced. Does Darwin's reading order reduce the surprise relative to the publication order, or does it increase it?

All results are relative to a null reading model that holds Darwin's original reading dates fixed and re-samples without replacement from his original reading list. The title selection at each reading date is constrained to those titles published before that date. In contrast to a null that includes permutations that neglect publication date, this restricted null captures the dynamics of publication in which a new work can unexpectedly change the information space. We are unable to resolve publication dates to less than a year; this occasionally implies an ambiguity in creation date for texts in the corpus that have the same year of publication. To solve this, we average our results over all possible within-year orders.

We compare the production and consumption of texts by considering only the texts Darwin recorded himself as reading. This allows us to make a direct comparison between the average surprise of the publication order and reading order within that set. But the production of these texts, of course, occurs in a much larger context, and it is reasonable to ask about the books that Darwin considered reading but did not, or an even more complete representation of the state of Victorian science constructed using all the scientific books available to him in Kent and London during these years. Here we focus on the books he chose to read, and the space he actually explored, partly for practical reasons and partly because we are interested in the decisions he made to order the texts, rather than the decision of whether to read them at all.

2.6. Bayesian epoch estimation

In the foraging literature, individuals are often assumed to persist in sustained periods of either exploration or exploitation. We call this an *epoch*. We are particularly interested in whether or not these epochs align with important events in Darwin's life. By using Bayesian models to infer epoch breaks, we can determine whether the data support a qualitative interpretation of the quantitative model.

Bayesian epoch estimation (BEE) models an epoch as a Gaussian distribution of relative surprise, in either the text-to-text or text-to-past case, with fixed mean and variance. Each epoch is defined by a beginning point (which is also either the end of the previous epoch, or the start of the data), an average level of surprise, and the variance around that average. For each time-series, the model then contains $3n - 1$ parameters, where n is the number of epochs.

Epoch switches are independently selected for the text-to-text and text-to-past measures. Each transition can then be interpreted as a change in Darwin's exploration and exploitation behavior at the local or global level. When a new epoch has higher average surprise than the one before, for example, we can understand Darwin as moving to a more exploration-based strategy.

Our one externally-set parameter is the total number of epochs, n . As n rises, it becomes easier and easier to fit the data; at some

point, we encounter the over-fitting problem, and in the extreme case we have as many parameters as we have data points. For each choice of n , our model returns a likelihood: the probability that the observed data were generated by the (best fitting) choice of parameters for that model. As n rises, so does the log-likelihood.

To determine the best-fit number of epochs, we use a simple model-complexity penalty, Akaike Information Criterion (AIC) (Akaike, 1974), to verify that the selected model is preferred, despite the addition of new parameters. The AIC penalizes a model's log-likelihood by the total number of parameters in the model; adding more epochs, in other words, must justify itself by a sufficiently large increase in goodness of fit. The AIC can be understood as an information-theoretic and Bayesian version of the chi-squared test, which attempts to maximize a model's predictive power (Burnham & Anderson, 2003).

To fix attention on the longest timescales in Darwin's life, we set the minimum epoch length to five years. See Appendix E for further discussion of the BEE model, for the likelihood space of epoch breaks for Darwin's readings, and for our AIC analysis.

3. Results

3.1. Exploration and exploitation

Over the 647 records in our corpus, Darwin's reading order led to a below-null average surprise, where the null is the average surprise of 1000 permutations of Darwin's reading order, constrained by each book's publication date (see Section 2.4).

On average, the KL divergence from text to text in the corpus is 10.78 bits compared to a null expectation of 11.41 bits ($p \ll 10^{-3}$). Meanwhile, Darwin's text-to-past average surprise is 2.96 bits in the data versus 2.98 bits in the null ($p = 0.02$). Darwin's average surprise, in both text-to-text and text-to-past, is lower than expected from a null model. While our rejection of this simple null model provides little new insight into the cognitive process of a reader's decision-making—which we expect to have some correlation from book to book—it is a crucial test of the sensitivity of our methods themselves.

A surprise-minimizing path is one that orders the texts so as to minimize the total sum of text-to-text, or text-to-past, surprise. Finding this shortest path amounts to a variant of the "traveling salesman problem", which is famously difficult to solve (reviewed in Cook (2011)). The greedy shortest path algorithm attempts to approximate the surprise-minimizing path by starting with the first text that Darwin read, and choosing as the next one to read the one with smallest KL divergence from the first, and so on, minimizing at each step either the text-to-text, or text-to-past KL depending on which quantity one is interested in.

While Darwin's path is lower in surprise than the null, it is far larger than many paths that can be found: the greedy shortest-path algorithm, for example, can reduce the text-to-text average surprise to 2.11 bits and text-to-past average surprise to 2.97 bits. Table 2 shows the raw local text-to-text and global text-to-past KL divergence data, along with the greedy shortest path single-visit traversals of the KL distance matrix.

The null, actual, and greedy shortest-path results show that Darwin has a focused reading strategy despite not following a pattern of pure surprise-minimization. Interestingly, the greedy shortest path is slightly longer than the path Darwin took in the global measure. This highlights how preference for exploration at the local scale—*i.e.*, not taking the closest book in topic space at each step—can lead to an unexpectedly efficient path in the global measure. See Appendix D for an alternative rank-order analysis reinforcing this null-model testing.

Table 2

Exploration habitsf. Average text-to-text (local) and text-to-past (global) KL Divergence (bits/step) over the reading path. Text-to-past KL is lower, as Darwin's reading spreads out to cover the topic space and lowers the information-theoretic surprise of subsequent books. Darwin's reading strategy is simultaneously more exploitative than would be expected of a random reader while also not following a strategy of pure surprise-minimization. Note that Darwin's order displays lower global surprise than the greedy shortest path, which demonstrates that selecting the next most similar book is not the best overall strategy for minimizing average global surprise over time.

	Local text-to-text (bits/step)	Global text-to-past (bits/step)
Darwin's order	10.78	2.96
Null (1000 permutations)	11.41 ± 0.28	2.98 ^{+0.04} _{-0.02}
(<i>p</i> -value)	≪ 10 ⁻³	0.02
Greedy shortest path	2.11	2.97

3.2. Readings over time

While Darwin is on average more exploitative, this is not necessarily true at any particular reading date. Darwin's surprise accumulates at different rates depending on time, as can be seen in Fig. 1 for the text-to-text case (top panel) and the text-to-past case (bottom panel). These figures plot the cumulative surprise relative to the null, so that a negative (downward) slope indicates reading decisions by Darwin that produce below-null instantaneous surprise (exploitation). Conversely, a positive (upward) slope indicates decisions that are more surprising than the null (exploration).

Over the entire corpus, as we know from the previous section, Darwin's cumulative surprise is below the null expectation, showing an overall bias towards both local and global exploitation. Tracking the slopes in these charts over time, however, allows us to see how Darwin moves between low-surprise and high-surprise choices on a range of timescales. The interaction of these decision rules at the text-to-text and text-to-past levels characterize Darwin's behavior.

3.3. Individual and collective

While many studies see scientific innovations as following large-scale cultural trends (Sun, Kaur, Milojevic, Flammini, & Menczer, 2013), individuals can also be understood as ahead of their time, pursuing connections and ideas before they are recognized by the culture as a whole (Johnson, 2010; Bliss, Peirson, Painter, & Laubichler, 2014). By ordering Darwin's readings by publication date, rather than reading date, we see how the culture gradually accumulates and assimilates content. We then compare how the culture produced these texts to how Darwin, in his reading, consumed them.

Whether surprise is higher in the reading order or the publication order is a substantive empirical question. There are good reasons to imagine that the reading order will be lower in text-to-text and text-to-past surprise than the publication order—but equally plausible reasons for the opposite.

Consider, first, the idea that the reading order surprise is the lower of the two. Informally, this would suggest that reader has been at least partially successful at reordering the large number of texts into a more systematic sequence, finding similarities between chronologically distant texts and making sense of a disordered, distributed discovery process. A reader like Darwin, particularly early in his career when benefiting from more senior mentors, might be expected to attempt to do something similar.

However, equally reasonable arguments work for the opposite direction, where reading order is more surprising than the

publication order. In this case, the reader is “remixing” the products of the past, sampling texts from different periods in such a way as to juxtapose thematically distant readings. While society accumulates themes gradually, in an exploitation regime, the reader is in an exploration regime, bringing them into unexpected contact.

While both possibilities appear *a priori* plausible, the data decisively favor the second. Fig. 2 shows the text-to-text and text-to-past cumulative surprise for Darwin's reading order (solid line) compared to the publication date order (dashed line). Since volumes are published and read at different times, the *x*-axis is now ordinal (*i.e.*, by position in the reading or publication sequence), rather than temporal (*i.e.*, by date read or published). This allows us to compare his reading order to the publication order independent of time.

Compared to Darwin's reading practices, cultural production has lower rates of surprise. While cumulative text-to-text surprise for Darwin often shows either flat or positive (above-null text-to-text surprise) slope, the publication order path is less explorative in both text-to-text and text-to-past cases. These findings, both at high levels of statistical significance ($p \ll 10^{-3}$), provide strong evidence for the remixing hypothesis.

3.4. Strategy shifts between biographically significant epochs

Between 1837 and 1860, Darwin's three major intellectual projects are reflected in his publication history. First, he began assembling his research journals on the geology and zoology from the voyage of the *HMS Beagle*. The last of these nine volumes was published in 1846. A second epoch can be dated from 1 October 1846 when, while assembling the last of his *Beagle* notes, Darwin discovered a gap in the taxonomic literature concerning the living and fossil *cirripedia* (or barnacles) (Darwin, 1838-1851a). After a period of intense work, he published four volumes on the taxon from 1851 to 1854. A final epoch begins with his journal entry on 9 September 1854, marking the day he began sorting his notes for a major work on species (Darwin, 1838-1851b). The revolutionary *Origin of Species* was published on 24 November 1859. These dates define three intellectual epochs: (1) from the beginning of records in 1837 to 30 September 1846; (2) from 1 October 1846 to 8 September 1854; and (3) from 9 September 1854 to the end of records in 1860.

We use the text-to-text and text-to-past models to characterize the exploration and exploitation of Darwin's reading behavior in each epoch. In instances where Darwin's average KL-divergence is above the null (more positive), Darwin is more exploratory. In instances where Darwin's average KL-divergence is below the null (more negative), Darwin is more exploitative. The degree to which he is in either mode is shown by the magnitude of the number. Table 3 shows these values.

Darwin's three biographical epochs are characterized by major shifts in both text-to-text and text-to-past surprise. Darwin begins, in epoch one, in an exploitation mode in both text-to-text and text-to-past. His turn to the barnacles in 1846 is marked by a shift from exploitation towards exploration at the text-to-past level (global shift to new area), and an intensification of his exploitation strategy at the local, text-to-text level (increased focus in this new area). In the third epoch, when Darwin “began sorting his notes for Species Theory” (Darwin, 1838-1851b), text-to-past remains in the exploration mode; text-to-text now shifts to exploration as well.

3.5. Unsupervised detection of strategy shifts

In addition to using Darwin's personally-specified epochs, we use a Bayesian model (Bayesian Epoch Estimation [BEE], see Methods) to estimate epoch breaks from text-to-text and text-to-past




Fig. 1. Epochs of exploration and exploitation in Darwin's reading choices. Text-to-text (top) and text-to-past (bottom) cumulative surprise over the reading path, in bits. More negative (downward) slope indicates lower surprise (exploitation); more positive (upward) slope indicates greater surprise (exploration). The three epochs, identified by an unsupervised Bayesian model, are marked as alternating shaded regions with key biographical events marked as dashed lines and labeled in the top graph. The first epoch shows global and local exploitation (lower surprise). The second epoch shows local exploitation and global exploration (increased surprise, in text-to-past only). The third epoch shows local and global exploration (higher surprise in both cases).

surprise alone. This process determines inflection points for Darwin's behavior without reference to outside biographical facts, allowing us to determine the extent to which the intellectual epochs identified by traditional, qualitative scholarship align with purely information-theoretic features of his reading.

For text-to-text surprise, we find the boundary at 16 September 1854 (log-likelihood relative to not having a boundary: $\Delta\mathcal{L} = 2.61$; 14 times as likely as not having a boundary). This is within 1 week of his journal entry on 9 September 1854 marking the start of his synthesis. For text-to-past surprise, we find the boundary at 27 May 1846 ($\Delta\mathcal{L} = 6.17$; 479 times as likely). The difference between the automatically-selected date and his recorded start date on 1 October 1846 suggests the need for further investigation of what our results suggest was a period of relative uncertainty for Darwin about his research plans.

The exploration-exploitation characteristics of these epochs are shown in Table 4. The close proximity of these automatically-detected breaks and the biographically significant epochs of the previous section confirm the central role of information-theoretic surprise in tracing the evolution of Darwin's search strategies.

Both the text-to-text and text-to-past models make highly simplifying assumptions about the nature of Darwin's reading. The text-to-text case makes the most severe assumption of all: that

Darwin's reading choices are conditional solely on the book just read. If Darwin's reading choices are strongly influenced by longer term memory (as seems likely), and it is these patterns which define the true epoch boundaries, it is natural that evidence for epoch boundaries in the text-to-text BEE model is weaker than the text-to-past case. In addition, our BEE makes the simplifying assumption that successive surprise values are independent draws from the distribution associated with that epoch.

4. Discussion

Models of cultural change often understand innovation as a multi-level combinatoric process, in which bundles of ideas are subject to cultural processes analogous to natural selection (Jacob, 1977; Wagner & Rosen, 2014). These evolutionary analogies typically consider change at the population level, as new ideas are created, spread, and modified by the crowd. A variety of recent studies covering conceptual formation in science, technology, and the humanities have taken this population-level perspective, including work on the recombination of patents (Youn et al., 2015), novelties (Tria et al., 2014), and citations (Garfield, 1979). Sociological studies of scientific practice have investigated how




Fig. 2. Darwin's reading order more exploratory than the culture's production. Text-to-text (top) and text-to-past (bottom) cumulative surprise over the reading order (solid) and the publication order (dashed). More negative (downward) slope indicates lower surprise (exploitation); more positive (upward) slope indicates greater surprise (exploration). In both cases, Darwin's cumulative surprise is higher than the publication order; in the second case, very significantly so. We mark the positions of three biographically significant books: Charles Lyell's *Principles of Geology* (3rd ed., 1837; read in 1837), Thomas Malthus's *An Essay on the Principle of Population* (1803; read on October 3, 1838), and Robert Chambers's *Vestiges of the Natural History of Creation* (1844; read on November 20, 1844). Darwin's juxtaposition of Lyell and Malthus, for example, is characteristic of how Darwin's reading strategies reordered the products of his culture.

Table 3

Information-theoretic correlates of biographically significant events. In this first table, we measure the relative surprise of Darwin's reading order by reference to dates derived from qualitative biographical work. The first major epoch of Darwin's intellectual life identified in this fashion corresponds to his post-*Beagle* work, when his readings were mostly in natural history and geology. Both text-to-text and text-to-past surprise remain low—a regime of simultaneous local and global exploitation. The second epoch, when Darwin turns to a study of barnacles, shows an increase in text-to-past surprise (new topics; exploration) coupled with a decrease in text-to-text surprise (smaller jumps within these new topics; exploitation). The third epoch, when Darwin begins to collect his notes for his “Species Theory”, is characterized by a rise in both text-to-text and text-to-past surprise. Now Darwin is neither repeatedly returning to well-covered topics (as in epoch one), nor turning his attention to a new, but narrow, range (as in epoch two), but rather ranging widely over new, previously understudied topics.

Start date	<i>Beagle</i> writings 2 October 1836	Barnacles 1 October 1846	Synthesis 9 September 1854
Text-to-text	-0.68	-0.96	0.32
Text-to-past	-0.09	-0.06	0.26

Table 4

Biographically significant events are detectable by unsupervised learning. Even in the absence of qualitative information about Darwin's life, our Bayesian model, based only on text-to-text and text-to-past surprise measurements finds—with only slight differences—the three historically-noted epochs of Table 3: (1) from the start of our records in 1837 until text-to-past surprise changes from exploitation to exploration in Spring 1846, (2) from Spring 1846 until text-to-text surprise changes from exploitation to exploration in Autumn 1856, and (3) from Autumn 1856 to the end of our data, when both (local) text-to-text and (global) text-to-past selection behaviors are in the exploration state. The automatically-selected and biographical epochs agree on these characterizations, with small variance in the second epoch's start date.

Start date	<i>Beagle</i> writings 2 October 1836	Barnacles 27 May 1846	Synthesis 16 September 1854
Text-to-text	-0.78	-0.76	0.21
Text-to-past	-0.11	-0.02	0.24

disciplines (Sun et al., 2013) or “communities of practice” (Bettencourt & Kaiser, 2015) are formed.

The mechanisms driving cultural innovation at the population-level cannot, however, be fully understood without taking into

account the cognitive processes that operate at the level of individual scientists. We have taken a step towards modeling these individual-level processes by studying the information foraging behavior of one preeminent scientist, using an information-theoretic framework applied to probabilistic topic models of his reading behavior. The information-theoretic measure we use to quantify surprise, KL divergence, connects both analytically and

empirically to linguistics (Calamaro & Jarosz, 2015; Hale, 2001; Levy, 2008; Light & Greiff, 2002; Martin et al., 2013; Resnik, 1993) and visual search (Demberg & Keller, 2008; Itti & Baldi, 2009). The LDA topic models that generate this information space also have cognitive correlates (Griffiths et al., 2007).

Our methods allow us to zoom in on Darwin's individual-level process to identify major epochs in his reading strategies. Over time, Darwin shifts towards increasing exploration as he prepares to write the *Origin*. Interestingly, the overall order we find empirically—exploitation then exploration—is in contrast to many of predictions derived from mathematical accounts of how optimal agents navigate the exploration–exploitation dilemma (Gittins, 1979). These generally predict that individuals begin with exploration (see, e.g., Berger-Tal et al. (2014)), shifting later to exploitation as they gain information about the environment.

Our results may be consistent with the standard models, if Darwin's early exploitation was guided by an earlier exploration phase prior to 1837. Or, it may be the case that the early phase of exploitation was necessary for Darwin to gain sufficient abilities or confidence to explore in a reliable fashion later. Finally, it is worth noting that some evidence in favor of these standard models can still be found in the short-term switch towards greater exploitation at the text-to-text surprise from the first to the second epoch. This suggests that these models may be useful at shorter timescales in an individual's life.

Our use of Charles Darwin allows us to validate our methods by reference to the extensive qualitative literature on his intellectual life. Having presented a general information-theoretic framework for describing the exploitation–exploration trade-off, we can now look at other searchers to see if other strategies exist for managing the exploitation–exploration trade-offs. Expanding these results beyond the Darwin test case will be essential to providing new empirical constraints on theories of how individuals explore the cultures of their time.

Reading records exist not only for elites, like Charles Darwin, but also for the general public. The maintenance of personal reading diaries is widespread, as evidenced by the more than 30,000 records in the UK Reading Experience Database (1450–1945),¹ and the 50 million registered users of Goodreads.² Further studies will be accelerated by advances in information retrieval techniques to find the full text for each entry in these reading diaries. These large-scale surveys can be complemented by in-depth studies of the “commonplace books” left by historical figures. These books record quotes, readings, and interactions that may become useful in their later intellectual life, and include Marcus Aurelius (Aurelius, 2012), Francis Bacon (Bacon, 1883), John Locke (Locke, 1706), and Thomas Jefferson (Wilson, 2014).

Our method also allows us to compare the individual and the collective. We have found, in particular, that Darwin followed a path through the texts that was more exploratory than the order in which the culture produced them. Our work reveals an important distinction between these two levels of analysis; underneath gradual cultural changes are the long leaps and exploration comprising an individual's consumption, combination, and synthesis.

Our cognitive analysis of these records builds upon decades of archival scholarship and innovations in the digital humanities. Darwin's industry extends beyond the bounds of the data we use here, however. During the *Beagle* voyage, he kept a library of 180–275 titles (Burkhardt & Smith, 1989). His retirement library contains 1484 titles (Rutherford, 1908). Darwin's handwritten marginalia in 743 of these books is currently being digitized by the Biodiversity Heritage Library. This retirement library contains

many texts not included in our study of his reading records, which only last through 1860. Finally, an extensive network of correspondents also contributed to Darwin's knowledge. The Darwin Correspondence Project³ contains over 15,000 letters to and from Darwin before 1869. A complete understanding of his information foraging will necessarily seek to understand this separate social process.

Darwin's sustained engagement with the products of his culture is remarkable. He averaged one book every ten days for twenty-three years, including works of fiction and foreign-language texts which are not part of the present analysis. For some months in our data, Darwin appears to be reading one book every two days, a fact even he was astonished by:

When I see the list of books of all kinds which I read and abstracted, including whole series of Journals and Transactions, I am surprised at my industry.
— *Autobiography of Charles Darwin* (Darwin, 1887, p. 119).

Darwin not only consumed information, it consumed him. In the words of Herbert Simon: “what information consumes is rather obvious: it consumes the attention of its recipients” (1991). Even the most ambitious individuals must confront and manage the limits of their own biology in allocating attention.

Standard theories for how individuals balance the exploration–exploitation tradeoff draw on classic work in the statistical sciences (Gittins, 1979) and machine-learning (Thrun, 1992), and often focus on determining mathematically optimal strategies for different environments (Cohen et al., 2007). Our work provides both new tools for the study of how individuals in the real world approach these problems, and new results on an exemplar individual.

5. Conclusion

Charles Darwin's well-documented reading choices show evidence of both exploration and exploitation of the products of his culture. Rather than follow a pure surprise-minimization strategy, Darwin moves from exploitation to exploration, at both the local and global level, in ways that correlate with biographically-significant intellectual epochs in his career. These switches can be detected with a simple unsupervised Bayesian model. Darwin's path through the books he read is significantly more exploratory than the culture's production of them.

To what extent the patterns we identify in Darwin, and his relationship to culture as a whole, hold for other scientists in other eras is an open question. The development of an individual is in part the history of what they choose to read, and it is natural to ask what patterns these choices have in common.

The methods we have developed and tested here represent the first application of topic modeling and cognitively-validated measures, such as KL divergence, to a single individual. These domain general methods can be used to study the information foraging patterns of any individual for whom appropriate records exist, and to look for common patterns across both time and culture.

Acknowledgments

We thank Peter M. Todd for extensive comments on a draft of this manuscript, as well as David Kaiser, Caitlin Fausey, Vanessa Ferdinand, and Jim Lennox for their feedback on presentations of this work. We also thank the audiences at various locations where the ideas in this paper were presented, especially within the Indiana University Cognitive Science Program. JM and SD thank the

¹ <http://www.open.ac.uk/Arts/reading/>.

² <https://www.goodreads.com/about/us> (accessed 2016 July 14).

³ <https://www.darwinproject.ac.uk/>.

Santa Fe Institute for their hospitality while this work was completed. We thank Tom Murphy for assistance with corpus curation and Robert Rose for programming assistance. Tools for corpus preparation and modeling were produced by Robert Rose and JM while supported by the National Endowment for the Humanities Digging Into Data Challenge (NEH HJ-50092-12, CA, co-PI). JM and CA were supported by an Indiana University (IU) Office of the Vice Provost for Research (OVPR) Faculty Research Support Program (FRSP) Seed Funding Grant and Bridge Funding Grant. JM was also supported by an IU Cognitive Science Program Supplemental Research Fellowship. Finally, we thank the anonymous reviewers and the editor for their constructive comments and honest inquiries.

Appendix A. Corpus characterization

Detailed characterization of Darwin's reading corpus, preparation methods, and software.

Appendix B. An introduction to Kullback-Liebler Divergence

An expanded, pedagogical introduction to Kullback-Liebler Divergence presenting two distinct interpretations of the measure.

Appendix C. Null model justification

Further justification of the null model, noting that any representation of Victorian science will be impoverished.

Appendix D. Text-to-text and text-to-past KL

Detailed analysis of the greedy shortest-path through Darwin's texts, showing that he indeed does not follow a surprise-minimization strategy. Alternate rank-order analysis of Darwin's reading order, indicating that while Darwin does not follow pure surprise-minimization, he does select the nearest neighbor more often than chance.

Appendix E. Bayesian epoch estimation

Further detail of Bayesian epoch estimation and further justify the independent selection of epoch break points.

Appendix F. Epoch model selection

AIC analysis for $k = 80$ topics, then repeated for 3 alternative models with $k = 20, 40, 60$ topics.

Appendix G. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.cognition.2016.11.012>.

References

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19, 716–723. <http://dx.doi.org/10.1109/TAC.1974.1100705>.
- Aurelius, M. (2012). *Meditations*. Mineola, NY, USA: Dover Publications. Translated by George Long.
- Azoulay-Schwartz, R., Kraus, S., & Wilkenfeld, J. (2004). Exploitation vs. exploration: Choosing a supplier in an environment of incomplete information. *Decision Support Systems*, 38, 1–18. [http://dx.doi.org/10.1016/S0167-9236\(03\)00061-7](http://dx.doi.org/10.1016/S0167-9236(03)00061-7).
- Bacon, F. (1883). *The promus of formularies and elegancies*. Longman, Greens and Company.
- Berger-Tal, O., Nathan, J., Meron, E., & Saltz, D. (2014). The exploration-exploitation dilemma: A multidisciplinary framework. *PLoS ONE*, 9. <http://dx.doi.org/10.1371/journal.pone.0095693>.
- Berra, T. M. (2009). *Charles Darwin: The concise story of an extraordinary man*. Baltimore, MD: John Hopkins University Press.
- Bettencourt, L.M., & Kaiser, D.I. (2015). *Formation of scientific fields as a universal topological transition*. Available from arXiv:1504.00319v1. SFI working paper 2015-03-009.
- Blei, D. (2012a). Topic modeling and digital humanities. *Journal of Digital Humanities*, 2, 8–11.
- Blei, D. M. (2012b). Probabilistic topic models. *Communications of the ACM*, 55, 77–84. <http://dx.doi.org/10.1145/2133806.2133826>.
- Blei, D. M., & Lafferty, J. D. (2007). A correlated topic model of science. *The Annals of Applied Statistics*, 17–35. <http://dx.doi.org/10.1214/07-AOAS114>.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *The Journal of Machine Learning Research*, 3, 993–1022.
- Bliss, N. T., Peirson, B. E., Painter, D., & Laubichler, M. D. (2014). Anomalous subgraph detection in publication networks: Leveraging truth. In *48th Asilomar conference on signals, systems and computers* (pp. 2005–2009). <http://dx.doi.org/10.1109/ACSSC.2014.7094823>.
- Burkhardt, F., & Smith, S. (1989). Books on the beagle. In *Darwin correspondence project*. <<https://www.darwinproject.ac.uk/books-on-the-beagle>>.
- Burnham, K. P., & Anderson, D. R. (2003). *Model selection and multimodel inference: A practical information-theoretic approach*. New York, NY, USA: Springer Science & Business Media.
- Calamaro, S., & Jarosz, G. (2015). Learning general phonological rules from distributional information: A computational model. *Cognitive Science*, 39, 647–666. <http://dx.doi.org/10.1111/cogs.12167>.
- Chun, M. M., & Wolfe, J. M. (1996). Just say no: How are visual searches terminated when there is no target present? *Cognitive Psychology*, 30, 39–78. <http://dx.doi.org/10.1006/cogp.1996.0002>.
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I Stay or Should I Go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362, 933–942. <http://dx.doi.org/10.1098/rstb.2007.2098>.
- Cohen Priva, U., & Austerweil, J. L. (2015). Analyzing the history of cognition using topic models. *Cognition*, 135, 4–9. <http://dx.doi.org/10.1016/j.cognition.2014.11.006>.
- Cook, W. (2011). In *Pursuit of the traveling salesman: Mathematics at the limits of computation*. Princeton University Press.
- Darwin, C. (1838-1851a). Darwin's 'Journal' (1809–1881). <<http://darwin-online.org.uk/content/frameset?viewtype=text&itemID=CUL-CUL-DAR158.1-76-&pageseq=62>> CUL-DAR158. Accessed on: 4 December 2015.
- Darwin, C. (1838-1851b). Darwin's 'Journal' (1809–1881). <<http://darwin-online.org.uk/content/frameset?viewtype=text&itemID=CUL-CUL-DAR158.1-76-&pageseq=46>> CUL-DAR158. Accessed on: 4 December 2015.
- Darwin, C. (1887). *The life and letters of Charles Darwin, including an autobiographical chapter*. John Murray.
- Demberg, V., & Keller, F. (2008). Data from eye-tracking corpora as evidence for theories of syntactic processing complexity. *Cognition*, 109, 193–210. <http://dx.doi.org/10.1016/j.cognition.2008.07.008>.
- Eliassen, S., Jørgensen, C., Mangel, M., & Giske, J. (2007). Exploration or exploitation: Life expectancy changes the value of learning in foraging strategies. *Oikos*, 116, 513–523. <http://dx.doi.org/10.1111/j.2006.0030-1299.15462.x>.
- Garfield, E. (1979). *Citation indexing – Its theory and application in science, technology, and humanities*. John Wiley & Sons, Inc.
- Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, 148–177.
- Goldstone, A., & Underwood, T. (2014). The quiet transformations of literary studies: What thirteen thousand scholars can tell us. *New Literary History*, 45, 359–384. <http://dx.doi.org/10.1353/nlh.2014.0025>.
- Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). *Topics in semantic representation*. <http://dx.doi.org/10.1037/0033-295X.114.2.211>.
- Griffiths, T. L., & Steyvers, M. (2004). Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101, 5228–5235. <http://dx.doi.org/10.1073/pnas.0307752101>.
- Gruber, H. E., & Barrett, P. H. (1974). *Darwin on man: A psychological study of scientific creativity*. EP Dutton.
- Hale, J. (2001). A probabilistic early parser as a psycholinguistic model. In *Proceedings of the second meeting of the North American chapter of the association for computational linguistics on language technologies* (pp. 1–8). Association for Computational Linguistics. <http://dx.doi.org/10.3115/1073336.1073357>.
- Hall, D., Jurafsky, D., & Manning, C. D. (2008). Studying the history of ideas using topic models. In *Proceedings of the conference on empirical methods in natural language processing EMNLP '08* (pp. 363–371). Stroudsburg, PA, USA: Association for Computational Linguistics. URL <<http://dl.acm.org/citation.cfm?id=1613763>>.
- Hills, T. T., & Adelman, J. S. (2015). Recent evolution of learnability in American English from 1800 to 2000. *Cognition*, 143, 87–92. <http://dx.doi.org/10.1016/j.cognition.2015.06.009>.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., & Couzin, I. D. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, 19, 46–54. <http://dx.doi.org/10.1016/j.tics.2014.10.004>.
- Hughes, J. M., Foti, N. J., Krakauer, D. C., & Rockmore, D. N. (2012). Quantitative patterns of stylistic influence in the evolution of literature. *Proceedings of the*

- National Academy of Sciences, 109, 7682–7686. <http://dx.doi.org/10.1073/pnas.1115407109>.
- Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Research*, 49, 1295–1306. <http://dx.doi.org/10.1016/j.visres.2008.09.007>.
- Jacob, F. (1977). Evolution and tinkering. *Science*, 196, 1161–1166. <http://dx.doi.org/10.1126/science.860134>.
- Jockers, M. (2013). *Macroanalysis: Digital methods and literary history. Topics in the digital humanities*. University of Illinois Press.
- Johnson, S. (2010). *Where good ideas come from: The natural history of innovation*. UK: Penguin.
- Klingenstein, S., Hitchcock, T., & DeDeo, S. (2014). The civilizing process in London's old bailey. *Proceedings of the National Academy of Sciences*, 111, 9419–9424. <http://dx.doi.org/10.1073/pnas.1405984111>.
- Kullback, S., & Leibler, R. (1951). On information and sufficiency. *Annals of Mathematical Statistics*, 22, 79–86. <http://dx.doi.org/10.1214/aoms/117729694>.
- Levy, R. (2008). Expectation-based syntactic comprehension. *Cognition*, 106, 1126–1177. <http://dx.doi.org/10.1016/j.cognition.2007.05.006>.
- Light, M., & Greiff, W. (2002). Statistical models for the induction and use of selectional preferences. *Cognitive Science*, 26, 269–281. http://dx.doi.org/10.1207/s15516709cog2603_4.
- Locke, J. (1706). *A new method of making common-place-books*. Printed for J. Greenwood, bookseller, at the end of Cornhil, next Stocks-Market.
- Macmillan, M. (2000). Restoring phineas gage: A 150th retrospective. *Journal of the History of the Neurosciences*, 9, 46–66. [http://dx.doi.org/10.1076/0964-704X\(200004\)9:1;1-2;FT046](http://dx.doi.org/10.1076/0964-704X(200004)9:1;1-2;FT046).
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2, 71–87. <http://dx.doi.org/10.1287/orsc.2.1.71>.
- Martin, A., Peperkamp, S., & Dupoux, E. (2013). Learning phonemes with a proto-lexicon. *Cognitive Science*, 37, 103–124. <http://dx.doi.org/10.1111/j.1551-6709.2012.01267.x>.
- Michel, J.-B., Shen, Y. K., Aiden, A. P., Veres, A., Gray, M. K., Team, T. G. B., ... Aiden, E. L. (2011). Quantitative analysis of culture using millions of digitized books. *Science*, 331, 176–182. <http://dx.doi.org/10.1126/science.1199644>.
- Mohr, J. W., & Bogdanov, P. (2013). Introduction – Topic models: What they are and why they matter. *Poetics*, 41, 545–569. <http://dx.doi.org/10.1016/j.poetic.2013.10.001>.
- Pirolli, P., & Card, S. (1999). Information foraging. *Psychological Review*, 106, 643–675. <http://dx.doi.org/10.1037/0033-295X.106.4.643>.
- Resnik, P. S. (1993). Selection and information: A class-based approach to lexical relationships. *IRCS technical reports series* (p. 200).
- Resnik, P. (1997). Selectional preference and sense disambiguation. In *Proceedings of the ACL SIGLEX workshop on tagging text with lexical semantics: Why, what, and how* (pp. 52–57). Washington, DC.
- Rutherford, H. W. (1908). *Catalogue of the library of Charles Darwin now in the Botany School*. Cambridge: Cambridge University Press. Compiled by H.W. Rutherford, of the University Library; with an introduction by Francis Darwin.
- Simon, H. A. (1971). Designing organizations for an information-rich world. *Computers, Communication, and the Public Interest*, 37, 40–41.
- Squire, L. R., & Zola-Morgan, J. T. (1991). The cognitive neuroscience of human memory since HM. *Annual Review of Neuroscience*, 34, 259–288. <http://dx.doi.org/10.1146/annurev-neuro-061010-113720>.
- Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton University Press.
- Sun, X., Kaur, J., Milojevic, S., Flammini, A., & Menczer, F. (2013). Social dynamics of science. *Scientific Reports*, 3. <http://dx.doi.org/10.1038/srep01069>.
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning*. Cambridge: MIT Press.
- Thrun, S. (1992). The role of exploration in learning control. In D. White & D. Sofge (Eds.), *Handbook for intelligent control: Neural, fuzzy and adaptive approaches*. Florence, Kentucky: Van Nostrand Reinhold.
- Todd, P. M., Hills, T. T., & Robbins, T. W. (Eds.). (2012). *Cognitive search: Evolution, algorithms, and the brain*. Cambridge, MA: MIT Press.
- Tria, F., Loreto, V., Servidio, V. D. P., & Strogatz, S. H. (2014). The dynamics of correlated novelties. *Scientific Reports*, 4. <http://dx.doi.org/10.1038/srep05890>.
- Underwood, T., & Sellers, J. (2012). The emergence of literary diction. *Journal of Digital Humanities*, 1.
- Uotila, J., Maula, M., Keil, T., & Zahra, S. A. (2009). Exploration, exploitation, and financial performance: Analysis of S&P 500 corporations. *Strategic Management Journal*, 30, 221–231. <http://dx.doi.org/10.1002/smj.738>.
- Van Hulle, D. (2014). *Modern manuscripts: The extended mind and creative undoing from Darwin to beckett and beyond. Historicizing modernism*. Bloomsbury Academic.
- Wagner, A., & Rosen, W. (2014). Spaces of the possible: Universal Darwinism and the wall between technological and biological innovation. *Journal of the Royal Society, Interface*, 11, 20131190. <http://dx.doi.org/10.1098/rsif.2013.1190>.
- Wilson, D. (2014). *Jefferson's literary commonplace book. Papers of Thomas Jefferson, Second*. Princeton University Press.
- Youn, H., Strumsky, D., Bettencourt, L. M., & Lobo, J. (2015). Invention as a combinatorial process: Evidence from US Patents. *Journal of the Royal Society*, 12, 1–8. <http://dx.doi.org/10.1098/rsif.2015.0272>.