

Surprise! Measuring Novelty as Expectation Violation

Jacob G. Foster*

Department of Sociology, UCLA, 264 Haines Hall, 375 Portola Plaza, Los Angeles, CA 90095

Feng Shi

Odum Institute for Research in Social Science, UNC Chapel Hill, 208 Raleigh St, Chapel Hill, NC 27517

James A. Evans*

Department of Sociology and Knowledge Lab, University of Chicago, 1126 E. 59th Street, Chicago, IL 60637; Santa Fe Institute, 1399 Hyde Park Road, Santa Fe NM 87501

Abstract

Novelty assessment is central to the study and management of innovation. Here we argue that new technologies, discoveries, and cultural products are deemed novel insofar as they seem unlikely or improbable, conditional on perceptions of prior knowledge and estimations of the inventive search process. This implies that novelty has different manifestations in fields with distinct prior knowledge and processes of invention; that measuring novelty is sensitive to context and therefore “objectively subjective.” Consequently, different novelty measures will be appropriate for different fields. We then survey and systematize existing *ex ante* novelty measures. We array them according to the speed of simulated search and the complexity of the space over which search is simulated. Using data from 157,595 U.S. patents granted in 2000 and 90,421 patents granted in 1990, we demonstrate that inventive fields vary in their distribution of novelty measures. We also find that familiar impact-based correlates of novelty are predicted by distinct characterizations of novelty in different fields. Consistent with our hypothesis that different inventive processes imply different novelty measures, we find that nearby fields, which share similar inventive processes, also manifest similar profiles in the relationship between novelty and impact. We conclude with principles of measure selection that could lead to more credible analyses of innovation.

Keywords: innovation, novelty, invention, science of science, technology, measurement

*Corresponding author

Email addresses: foster@soc.ucla.edu (Jacob G. Foster), jevans@uchicago.edu (James A. Evans)

1. Introduction

Novelty is a critical concept in the study of innovation. Research in innovation studies, the sociology of science, economics, economic sociology, management, and allied disciplines highlights the role of novel discoveries, technologies, and cultural products across many settings, e.g., scientific and technological advance, economic growth, and corporate strategy. The modern academy defines itself through novel contributions to human understanding. Entire economic sectors are devoted to the creation and exploitation of novel devices and designs.

Yet the measurement and modeling of novelty reveals a paradox. On the one hand, researchers—including ourselves—often invent new measurements rather than build upon previous ones, or select from among the broad field of measures with limited justification for one over the others. Uzzi et al. (2013) makes a pivotal contribution to our understanding of atypical combinations and scientific impact, but does so through introducing an atypical novelty measure. Cokol et al. (2005) introduces a novel, network-based measure of novelty with no reference to prior approaches, as do many others (Chen et al. 2009, Schilling and Green 2011, Leahey and Moody 2014). Novelty measures are extremely heterogeneous and sometimes correlate negatively, as we show below. On the other hand, such measures are typically introduced as transcending particular empirical contexts, implying broad or even universal identification of novelty. Both conditions cannot hold. Either each measure captures a distinctive but largely unspecified aspect of novelty, or all measures capture the same thing, which varies by inventive and evaluative context. In this essay, we will argue for and empirically demonstrate the latter.

Novelty measurement contradicts the pattern in many areas where social scientists have been systematic in developing measures for key constructs like income inequality (Cowell 2000) and residential segregation (Hutchens 2001). Invention, innovation, and discovery are multi-faceted, and defy simple, singular description. Yet the flexibility in measurement choice raises a temptation to cherry-pick the measure that produces the “best result”. Here we seek to identify and build connections between existing novelty measures by exploring the conceptual underpinnings of novelty assessments in the service of making novelty measurement more systematic, cumulative, accurate, and conceptually clear.

We first provide a formal theory of novelty production and perception as a phenomenon that varies by context (Section 2). Our theory is grounded in prevailing contemporary accounts of cognition, which place prediction at the heart of most cognitive processes (Clark 2015). Our central claim is that the production of novelty is grounded in the inventive process within a particular domain, and that novelty assessment derives from situated mental models of that process. What is deemed novel corresponds to what is surprising, given those models. There can be no universally objective novelty measurement because an object’s novelty is defined as new in relation to the perceived (or measured) qualities of other objects with which it is compared. But there is no object collection universally relevant across all times and domains. Novelty without context is a numerator without denominator or a difference without subtrahend. Moreover, perceptions and measurements are themselves collected with respect to anticipated evaluations, constraining our awareness of perceivable novelty.¹

Measures of novelty implicitly translate models of invention into formal measures². We review the many existing approaches to novelty measurement from this perspective (Section 3), placing them within a common conceptual framework by noting that these measures operationalize distinct ideas about (i) the way that inventors search and (ii) the complexity of the space in which they do so. Clearly stating the assumptions behind each class of measures suggests that different novelty measures will be most appropriate for different fields of invention.

We demonstrate this claim empirically in Section 4, using data from 90,421 U.S. utility patents granted in 1990 and 157,595 granted in 2000; with detailed results in the Appendix. We find that patents from different technology areas vary substantially in their distribution of novelty scores across measures. We

¹Was the asteroid believed responsible for the extinction of 75% of the earth’s plant and animal species 66 million years ago novel? To generations of dinosaurs, yes; to the solar system over billions of years, no.

²By the same token, judgments of novelty, such as those made by patent examiners when evaluating patent applications, are also conditioned on implicit mental models of invention.

further find that in different technology areas, distinct measures are important to the prediction of various dimensions of impact. Consistent with our claim that novelty measures capture perceived novelty when they identify features of the inventive process, we find that fields where patents most intensively cite one another—and presumably share inventive processes—have greater similarity in predictive novelty measures. This also allows us to characterize areas by the appropriateness of novelty measure. Finally, we compare predictive novelty measures across different US PTO technical centers, showing that the relationship between novelty measures and impact varies systematically across inventive domains. This analysis supports two methodological inferences. First, there is no hope for “one novelty measure to rule them all.” Hence any novelty measure should be approached with skepticism. For a given analysis, the choice of novelty measure should be justified by showing that it (a) maps onto domain-specific knowledge structures and inventive processes; (b) is most predictive of novelty correlates like impact; and/or (c) correctly tracks actor perceptions of novelty. Second, novelty measurement will become better and more useful insofar as we develop better formal representations of knowledge and inventive processes.

In the last part of the paper, we turn to this second methodological implication. In Section 5.1, we argue that enriching the space of features used to characterize inventions can bring novelty measures closer to reality, allowing analysts to trace evolution in the structure of existing knowledge and how this channels new exploration across it. We discuss the “black-boxing” process whereby frequent combinations collapse into single components, accelerating the creation of new, complex combinations. In Section 5.2, we return to the subjective character of novelty assessment, and show how to give rigorous, objective treatment to that subjectivity: Bayesian surprisal allows scholars to explicitly incorporate individual variation in knowledge, beliefs, and experience into measures of perceived novelty.

We hope our theorization of novelty and synthesis of novelty assessment will improve evaluation of novelty, invention, and innovation, checking the natural tendency to (inadvertently) cherry-pick measures that best tell a preferred story. Above all, we hope it will inspire renewed attention by innovation scholars to the measurement of novelty, drawing on insights from the psychology and cognitive science of creativity, invention, and novelty perception. Better novelty measurement will lead to better novelty management, with profound implications for technology, science, commerce, art, and culture.

2. What is novelty? How should we measure it?

Novelty is easy to define but hard to measure. The OED defines it as “the quality or state of being novel.” Novel is a bit more interesting: “young, fresh; newly made or created; of a new kind or nature; strange; previously unknown; original” (OED 2007). This definition draws our attention to two crucial facets of novelty: first, that it is bound up with processes of making or creating; and second, that its evaluation has a context-sensitive and subjective dimension—what is “previously unknown” or unfamiliar to one may not be to another.

Scholars of innovation often distinguish *ex post* measures, which make reference to events that happen after an invention appears, from *ex ante* measures, which refer only to events that lead up to its debut. *Ex post* measures like diffusion and impact are appropriate for measuring the success or failure of an invention, such as whether it becomes an *innovation* by spreading or altering the scientific or technical landscape (Benner and Tushman 2003, Christensen 2011, Tushman et al. 1986, Padgett and Powell 2012). Such measures implicitly or explicitly model the process by which an invention becomes *important* (Wang et al. 2013, Gerrish and Blei 2010). For examples of *ex post* measurement, consider the frequent use of patent and paper citations to identify critical innovations, or the disruption measure by Funk and Owen-Smith that identifies breakthrough technologies by identifying cases where a new patent eclipses citation attention to the contributions on which it built (Funk and Owen-Smith 2016, Wu et al. 2019, Lin et al. 2021)³. Note that such *ex post* measures cannot quantify the *potential* for an idea or technology to be influential

³An increasing trend is to use both *ex post* measures of innovation and *ex ante* measures of novelty by measuring the (un)predictability of future from past (i.e., novelty) and the (un)predictability of past from future (i.e., persistent innovation) (Kelly et al. under review, Barron et al. 2018).

or transformative until after it begins to diffuse and disrupt. Measuring what happens after an invention appears does not independently measure its novelty, but confounds this with whether that novelty becomes important.

By contrast, *ex ante* measures of novelty do not require observations of the subsequent career of an invention (Verhoeven et al. 2016). Unlike *ex post* measures, which capture how an invention is received, *ex ante* measures assess how its qualities relate to earlier work, and how it is *perceived* as a result. In characterizing the relationship between an invention and prior art or knowledge, *ex ante* measures, in turn, implicitly or explicitly make claims about and hence model the process through which an invention or a discovery is made.

We note another important class of novelty measures focus on novelty *intentions* (or *voluntate* measures if we insist on Latin approbatives), such as investments in or intensity of original research and development (Lee et al. 2015, Lee 2015, Cohen 2010, Romer 1986, 1994, Kemp and Pearson 2008). While these measures likely elucidate a critical source of novelty, the intention and resources required to explore techno-scientific novelty likely has a complex relationship with realized novelty that is beyond the scope of this article. In this paper, we focus on measures of *ex ante* novelty, not on investments in novelty and not *ex post* measures of diffusion, innovation and impact ⁴.

For the purposes of this paper, we pragmatically define discovery as synonymous with invention and the creation of cultural artifacts. New technologies, enterprises, theories, paintings, symphonies—all result from creative human effort and involve some novel deviation from prior work, which makes them “new” (North 2013). How do social actors recognize this new-ness? Consistent with contemporary theories of predictive coding, we propose that they judge a creation as novel when it violates expectations (Rohrmeier and Koelsch 2012, Clark 2015). The more expectations are violated, the more novel a creation seems. What sets expectations? Familiarity with prior art and beliefs about the process of invention. To take a musical example, the mature works of Arnold Schoenberg were perceived as radically new. They represented a massive break with prevailing compositional methods, violating expectations set by prior art and standard compositional process (Ross 2007). The works of Maurice Ravel, his contemporary, violate expectations less, and consequently seem less novel. Ravel’s work includes techniques outside the familiar inventive process (e.g., gamelan music), but far fewer such techniques than Schoenberg’s (Mawer 2000). Schoenberg was perceived as more novel than Ravel because his music was more improbable given the “normal” creative process⁵.

This experiential account of novelty can also be grounded in the anthropologist Alfred Gell’s work on art (Gell 1998) ⁶. Gell argued that our experience of works of art is mediated by an attempt to infer (via abduction) the generative processes that produced them. Art that “enchants” is precisely the art whose generative processes cannot be inferred. The analogy to invention, discovery, and other forms of creation is straightforward. The role of *abduction* in Gell’s account also helps to clarify the importance of context. Abductive inference depends *precisely* on the stock of familiar ideas and procedures out of which an explanatory process can be assembled. The ancients were astonished by the work of a legendary glassmaker who could produce “flexible glass;” as Isidore of Seville reports, because they could not fathom how he had done it, they killed him to stop the novelty from spreading. Contemporary manufacturers, familiar with recent techniques for producing fiber optic cables and Gorilla glass, would presumably have a less lethal response.

Formally, we claim that perceived novelty is a decreasing function g of the mentally simulated probability of a discovery or invention D , conditional on the prior knowledge or art K and the creative/search process S through which discoveries or inventions are typically conceived and achieved⁷:

⁴Like all dichotomies, the novelty/impact dichotomy is somewhat artificial. Novelty affects impact, and impact—over time—necessarily shapes perceptions of novelty.

⁵There are, of course, scope conditions here. Improbable outcomes may be perceived as “crazy” if they violate expectations so thoroughly that they are not decipherable under current practices. But this can change with time. For example, quantum biology has shifted from “crazy idea” to intriguing possibility over the past decade.

⁶We thank Christo Sims for bringing this work to our attention.

⁷The conditioning is necessary because the probability of a discovery depends on what is already known (K) and on the

$$Nov = g(P(D|K, S)), \frac{dg}{dP} \leq 0. \quad (1)$$

If the estimated probability of a discovery D is low, given prior knowledge K and inventive process S , then its perceived novelty is high: it is surprising⁸. If the probability is high, then perceived novelty is low: the invention was an obvious or expected extension of prior work. On this account, measures of novelty simultaneously assess (1) the degree of perceived difference between a new object and the population of pre-existing objects *and* (2) the perceived improbability of arriving at the new object, assuming that components from previous objects constitute most of the raw materials. In short, measures of novelty assess the surprise induced by a focal object.

This connection can be made even tighter using explicit formalizations of surprise. An early, information-theoretic formalization reinterpreted the notion of self-information as surprise. Myron Tribus renamed self-information “surprisal” to signal the surprise experienced by the receiver of a message or the observer of an experiment upon observing a particular outcome (Tribus 1961). If I is the self-information or surprisal associated with experiencing outcome x_i of random variable X , then $I(x_i) = -\log(P(X = x_i))$.

Surprise is a decreasing function of the probability of experiencing a given outcome from a random variable, such that high probability events are banal and low probability events are surprising. This approach has been used in measures of the surprise associated with scientific discovery (Foster et al. 2015) and the surprise experienced by viewers of movies and sports matches (Ely et al. 2015). Variations of this basic idea will show up in many of the measures discussed in 3; this is unsurprising, given that existing measurements of novelty can be interpreted as simulations of the process by which discoveries or inventions have been achieved. If novelty measures are based on models of discovery and invention, however, is there any *organization* to those models? We claim that there is.

3. Organizing Novelty Measures

As we noted in the introduction, a plethora of data representations and algorithms have been used to measure technological, scientific and cultural novelty. Many of these measures do not explicitly articulate a simulation of discovery. Nevertheless, all measures capture some distance or divergence between a product and its predecessors, i.e., between the discovery D and pre-existing knowledge K . We argue that each *implies* a model of the process S by which a new creation was invented. A first set of measures, which we call *categorical*, identifies novelty with the presence of any indisputably novel *component*. Categorical measures have a simple implicit model of discovery: novel inventions simply *pop* into the minds of their inventors at some rate. Other measures, which capture gradations of novelty, have richer models of the discovery process. These *complex* models vary along two primary dimensions shown in Figure 1: (1) the speed of simulated search underlying discovery and invention; and (2) the complexity of the space over which search is simulated. Note that these two dimensions track the two factors on which actors condition when estimating the probability of some invention: the nature of the search process S and the structure of the space of pre-existing knowledge K . In other words, our conceptualization of novelty spells out the two axes along which we expect novelty measures to vary. Of these two axes, speed of search S is the principle axis of variation. We consider three approaches positioned along this axis, from slow to fast: *configurational*, *structural*, and *combinatorial* novelty. We consider each of these approaches in turn, providing examples and describing the assumptions under which they might be expected to work.

3.1. Categorical

The simplest family of measures captures *categorical novelty*: the determination that something is new because it includes a new component or combination thereof. For example, consider patent counts, which

“search” process (S) involved in creation. Anton Webern’s music presumably struck his teacher Schoenberg as less novel than it did the average listener.

⁸Of the discovery of the muon, I.I. Rabi famously said, “Who ordered that?”

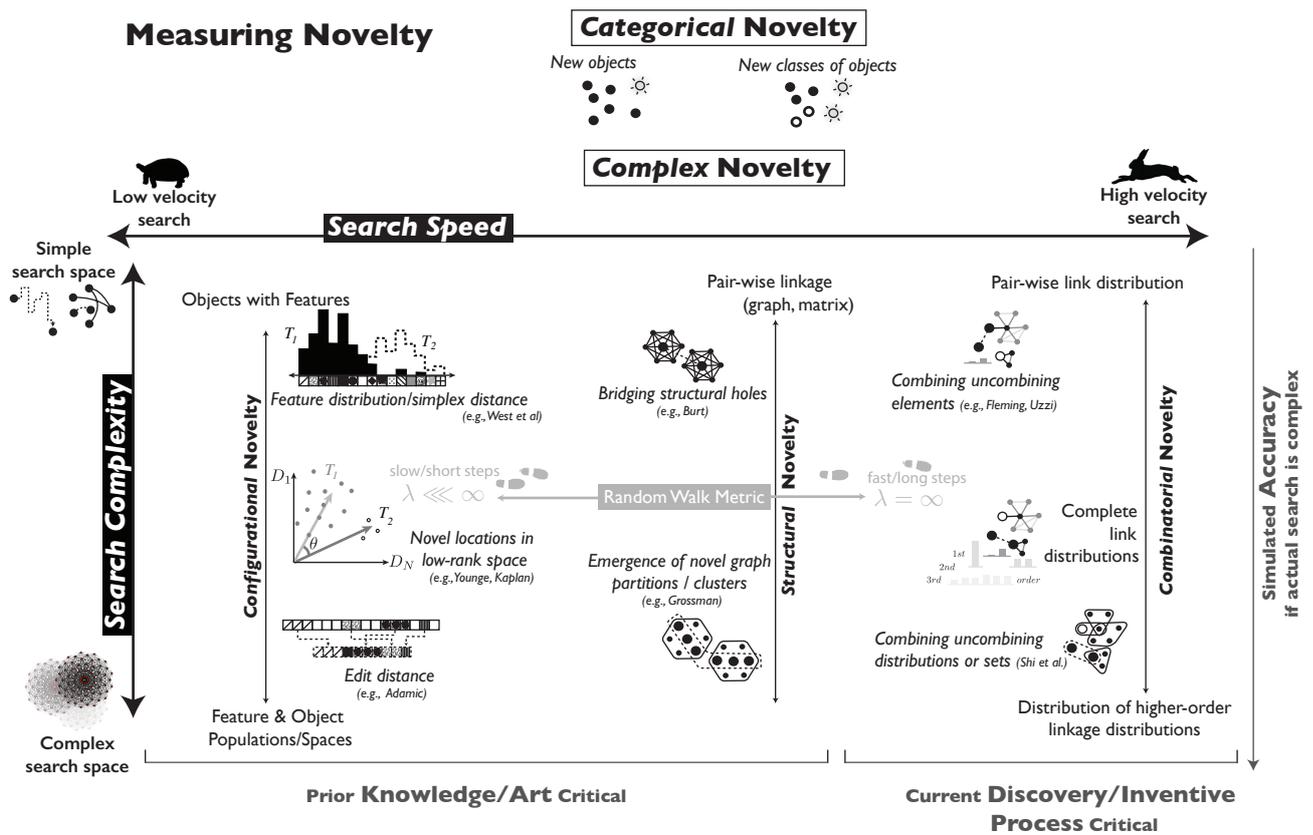


Figure 1: Novelty measures arrayed along the dimensions of (1) search speed and hence distance involved in typical search and (2) complexity of the search space.

represent technical combinations that have passed the innovation “novel” and “nonobvious” criteria required for their prosecution (Jaffe et al. 1993, Cohen 2010), counts of innovations identified through expert appraisal (Geroski and Pomroy 1990, Blundell et al. 1999), counts of new product announcements in trade journals and magazines (Coombs et al. 1996, Fosfuri and Giarratana 2009), clinical trials, counts of drug products, journal articles, or number of new molecular entities (Dubois et al. 2015), or worlds fairs and related exhibitions of new technology (Moser 2005). Other categorical novelty measures define novelty as a function of unprecedented content, such as Cokol et al. (2005) who defined novel biomedical research articles as those that introduced new chemicals. In their simplest form, assessments of categorical novelty are binary: an invention includes “something” new or it does not. Categorical assignments can be broadened to include multiple, graded categories of novelty. In Cokol et al. (2005) and Foster et al. (2015), articles that did not debut new chemicals but asserted new relationships between existing chemicals were scored as less novel, but nonetheless judged more novel than articles that introduced no new relationships at all. This simple tripartite novelty taxonomy lies atop a simple implicit simulation of discovery: that introducing a new chemical is more improbable than searching for a new chemical relationship, which is more improbable than tuning a known relationship⁹. These measures have been used in cases where the underlying inventive components are clear, well-defined, and easily identified.

This categorical approach runs roughshod over important detail. When creative products are opened up and their internal complexity examined, no component is ever completely new (North 2013). More sophisticated measures are needed to identify combinations of components that are *more or less* probable. We turn now to these *complex novelty* measures, describing them in order of inventive search speed implied. In other words, we array these measures along the S axis, corresponding to coarse-grained models of the search and discovery process.

3.2. Configurational Novelty

Configurational approaches to novelty assessment directly assess differences between present and past inventions in terms of their configuration of components. By explicitly using past knowledge and art as the baseline, these approaches assume that the inventive process reflects creativity in proportion to the distance traveled from that past. Configuration-based approaches represent inventions as sets of features and embed them in an “inventive space” spanned by these features. Discovery involves moving from established to novel positions in this space via shifts in feature values. Configurational approaches assume that search is typically slow, encoding a gradual, tinkering mode of invention, rather than radical recombination. Large movements in this space are highly improbable; they correspond to substantial departures from existing inventions, which were presumably used as template.

3.2.1. Feature Sets.

The most common configurational measures assess an object’s novelty relative to prior objects in a “bag of features” representation that ignores an invention’s internal structure. These measures build on the idea that each discrete feature of the invention (e.g., transistor, heat sink, word, image or musical motif) can be viewed as a dimension. An invention can then be characterized as a feature vector, registering a location within the high-dimensional feature space. In the simplest case, features are binary: present or absent. Inventions are then described by strings of *binary features*. For example, invention X is described as the feature string: $\mathbf{f}^X = f_1^X f_2^X \dots f_n^X$ of fixed length, where n is the total number of possible features and $f_i^X \in \{0, 1\}$. Inventions in this framework correspond to vertices on the hypercube of dimension n (Foster et al. 2013a) and distance between them is the Hamming distance. As a result, $Nov(\mathbf{f}^Y | \mathbf{f}^X) = \sum_i |f_i^X - f_i^Y|$, where the novelty of invention \mathbf{f}^Y is assessed in comparison to earlier invention \mathbf{f}^X . This distance could be averaged across several previous inventions, or calculated with respect to the *closest* previous invention.

More nuanced measures assess difference by weighting the features, perhaps as a function of the frequency with which each feature appears in past and present inventions. This weighting process still represents

⁹Note that measures that produce categorical variables are not necessarily categorical novelty measures as defined here; *all* the measures below can be converted into a formally categorical measure by establishing a cutoff, everything above which is novel and below which is not.

inventions as a *feature vector*. For example, Kelly et al. (under review) and Nanda et al. (2013) both use the cosine distance between document feature vectors to score the novelty of patents relative to the field of prior patents. Patent features comprised the full space of words in patent descriptions, scaling each patent’s feature values with term frequency-inverse document frequency weightings (*tf.idf*). Under this weighting scheme, more distinguishing features make a larger contribution to an invention’s position in the space of all possible inventions. This measure allowed Kelly and colleagues to improve the prediction of company valuations, just as Nanda, Younge and Fleming demonstrate how venture capital-backed start-ups in the renewable energy sector generate more novel patents than incumbent firms (Nanda et al. 2013). In practice, analysts commonly reduce the number of dimensions using matrix factorization approaches like singular value (or some other spectral) decomposition (Deerwester et al. 1990), latent variable methods like factor analysis or topic modeling (Blei et al. 2003), or neural networks like Google’s popular *word2vec* family of word embedding routines (Mikolov et al. 2013b,a) and recent alternatives (Pennington et al. 2014, Joulin et al. 2016, Devlin et al. 2018). A threshold-based variant of this approach using Bi-directional Encoder Representations from Transformers (BERT) was recently used to measure the novelty of research produced during the COVID crisis, identifying an increase in the average distance between paper elements (Liu et al. 2020). Specifically, Liu et al calculated the distance between bio-entities extracted from a paper within the BioBERT model (Lee et al. 2020), then counted the number of entity pairs in which their distance was in the 90th percentile of the pairwise distance distribution, normalized by the number of such pairs in the paper. Such approaches weight and merge similar features to generate reduced representations $\tilde{\mathbf{f}}^A, \tilde{\mathbf{f}}^B$ of the objects in question, on which familiar similarity measures can then be deployed as continuous measurements or through cutoffs. See Table 1 for stylized formalisms and a thumbnail characterization.

Feature vectors can also be treated as draws from an underlying (multinomial) distribution T^A . Novelty is operationalized as the difference between past T^A and present T^B probability distributions over features, using measures like the Kullback-Leibler divergence, Jensen-Shannon distance¹⁰, or Wasserstein / “earth-mover’s” distance. This approach is common in information retrieval tasks seeking to retrieve novel results against a repeated query (Larkey et al. 2002). As with the feature vector approach, analysts typically reduce the dimension of the space in which inventions are located as described above. The novelty of new inventions or documents can then be assessed based on their divergence or probability shift from older ones (Delort and Alfonseca 2012). Consider how Barron et al. (2018) used LDA topic modeling to factor words from 40,000 speeches delivered during the French Revolution’s first parliament into 100 topics (i.e., word distributions). They then identified novel speeches and speakers (e.g., the revolutionary proposals of the leftist Robespierre) as a function of their historical unpredictability, based on the KL divergence between the topic loadings for that speech and those prior: $Nov(\tilde{\mathbf{f}}^B | \tilde{\mathbf{f}}^A) = - \sum \tilde{\mathbf{f}}_i^B \log \left(\frac{\tilde{\mathbf{f}}_i^A}{\tilde{\mathbf{f}}_i^B} \right)$ (see Klingenstein et al. 2014, Murdock et al. 2015, for similar applications).

Feature set approaches are most useful when inventions can be well-described as collections of features, ignoring the detailed relationship between features. Distance is associated with novelty, and novelty with improbability. The most probable inventions are ones that make modest, incremental changes in a few features, relative to the pool of established inventions. Highly improbable inventions “travel far” from established invention, dramatically shifting the distribution of features. Alternatively, when inventions are viewed as draws from an underlying probability distribution of feature co-appearance, incremental invention involves a small tweak to that distribution—a little more of this, a little less of that—while low probability, high novelty invention involves significant change, such as from *Pride and Prejudice* (Austen 1813) to *Pride and Prejudice and Zombies* (Austen and Grahame-Smith 2009).

¹⁰The Kullback-Leibler divergence is commonly used to measure the “distance” between two distributions P and Q , but it is asymmetric, i.e., $KL(P, Q) \neq KL(Q, P)$, and it does not satisfy the triangle inequality. For this reason, it is not a true metric (technically, it is a premetric). The symmetrical Jensen-Shannon divergence *can* be used to define a metric, but effectively assesses the distance of each distribution to a midpoint.

3.2.2. Sequences and Complex Configurations

More complex configurational measures account not only for unstructured presence, absence, frequency or likelihood of features, but also their precise position in space and time. Consider an analysis of the evolution of quotes on social media, where Simmons, Adamic, and Adar (Simmons et al. 2011) calculate the edits required to transform an original quote \mathbf{s}^A into one being passed along \mathbf{s}^B . The number of insertions, deletions and substitutions needed to transform source string \mathbf{s}^A into string \mathbf{s}^B provides a measure of its relative novelty, $Nov(\mathbf{s}^B|\mathbf{s}^A)$. This measure deviates from a feature-based approach by detecting the difference between quotes containing the *same* components assembled in *distinct* sequences.

Economist Daniel Gross uses configuration-based methods not on one-dimensional sequences but two-dimensional (2D) raster images to investigate the relationship between competition and creativity in an online logo design challenge. Gross measures the novelty of an artist’s logo design as a function of its 2D difference from prior submissions, including their own (Gross 2014). He uses two algorithms to compute that difference: perceptual hash and difference hash, which perform a 2D edit-distance on the new logo image and its predecessors.

We note that in both examples, the precise internal configuration of new inventions is critical. Edits occur at *particular* positions along the string or in the image. As more fine-grained data on the internal structure of inventions becomes available, analysts can and should develop more detailed, configuration-based approaches to novelty measurement that take into account positions in three dimensional space or four dimensional space-time (e.g., when two components of a transistor are materially interconnected within a new radio-signaling device). These novelty measures would allow analysts to operationalize classic notions like architectural innovation (Henderson and Clark 1990) at scale. Such approaches make sense when inventions can be defined by a set of topological or geometric adjacencies among their components (e.g., component A is connected to or a certain distance from component B at time t) and when invention occurs through local alterations to that structure. In such a context, inventions with few shifts are more probable than those with many, because each edit incurs a cost of time, effort, creativity, etc. With the maturation of deep-learning approaches (e.g., Deep auto-encoders, Long Short-Term Memory neural networks, Generative Adversarial Networks, etc.) that infer a generative model capable of producing novel instances of inventive objects (Goodfellow et al. 2016) such as music, images, text, or video, novelty analysts can now make novelty assessments that take advantage of complete configurational information in their models of discovery.

3.2.3. Summary.

Configuration-based novelty measurements can be pursued with features of various kinds; with any low-rank description of the data (e.g., factors, spectra, clusters, partitions); and, increasingly, with detailed maps of internal configurations. For each configurational novelty measure, a distance or divergence in feature space is computed between the configuration of components in a new technical, scientific or cultural product and one or more predecessors. These measures correspond to a simulated inventive process in which a creator searches for new products by making incremental modifications to past ones. Longer searches require more effort, luck, or creative genius than short ones. Perceived novelty is then proportional to the configurational distance traversed by those simulated searches.

Category	Sub-category	Stylized Formalism / Instance(s)	Mental Model	Example
Configurational			Foreground: Inventions as <i>collections</i> of features.	
			Invention by replacing features.	
	<i>Binary features</i>	$Nov(\mathbf{f}^B \mathbf{f}^A) = \sum_i f_i^A - f_i^B $ Hamming distance	Features are present or absent. Invention by adding/subtracting features.	(Foster et al. 2013b)
	<i>Feature vector</i>	$Nov(\mathbf{f}^B \mathbf{f}^A) = 1 - \frac{\sum_i f_i^A f_i^B}{\ \mathbf{f}^A\ \ \mathbf{f}^B\ }$ Cosine distance	Features may be weighted (e.g., <i>tf.idf</i>). Invention by moving in feature-space.	(Kelly et al. under review)
	<i>Reduced vector</i>	$Nov(\tilde{\mathbf{f}}^B \tilde{\mathbf{f}}^A) = 1 - \frac{\sum_i \tilde{f}_i^A \tilde{f}_i^B}{\ \tilde{\mathbf{f}}^A\ \ \tilde{\mathbf{f}}^B\ }$ Cosine distance	Low-dimensional feature representation (e.g., PCA). Invention by moving in relevant feature-space.	(Pennington et al. 2014)
	<i>Distribution</i>	$Nov(T^B T^A) = \sum_i T_i^B \log \frac{T_i^B}{T_i^A}$ Kullback-Leibler divergence	Features drawn from distribution or mixture (T^A). Invention by moving in distribution or mixture-space.	(Murdock et al. 2015)
	<i>Sequence</i>	$Nov(\mathbf{s}^B \mathbf{s}^A) = \#$ local changes. Edit distance	Location of features matters (e.g., string, image). Invention by local editing.	(Simmons et al. 2011)
Structural			Foreground: Network <i>structure</i> of prior knowledge.	
			Invention by linking features in clusters.	
	<i>Bridging</i>	$Nov(i, j) = 1 - \delta(C_i, C_j)$ Categorical community distance	Frequently-linked components i form clusters C . Invention by combining features (easier within cluster).	(Foster et al. 2015)
	<i>Structural holes</i>	$Nov(i) \propto \sum_{st} \frac{n_{st}^i}{g_{st}}$ Betweenness centrality	Knowledge forms clusters; spanning is rare and difficult. Invention by combining features (most novel span clusters).	(Chen et al. 2009)
	<i>Emergent</i>	$Q = \frac{1}{2m} \sum_{ij} (A_{ij} - \frac{k_i k_j}{2m}) \delta(C_i, C_j)$ Modularity, map equation, etc.	New structures emerge from linking patterns. Novel contributions re-organize network structure.	(Rosvall and Bergstrom 2010)
Combinatorial			Foreground: Recombinant <i>process</i> of invention.	
			Invention by recombining features.	
	<i>Unipartite</i>	$Nov(i, j, \dots) = g(k_i, k_j, \dots)$ e.g., $\propto (\sum_i k_i)^{-1}$ or $\frac{1}{k_i k_j}$	Components i, j characterized by combinability k_i . Combining rarely-combined components more novel.	(Fleming 2001) (Uzzi et al. 2013)
	<i>Multipartite</i>	$Nov(i, j) = f[c_0 + \sum_\alpha c_\alpha P^\alpha(i, j)]$ Random walks on hypergraphs	Multiple types of component densely interwoven. Invention through more- and less-likely paths.	(Shi et al. 2015)

3.3. Structural Novelty

Structural approaches to novelty measurement place the relational structure of prior knowledge in the foreground. Instead of representing inventions as feature vectors in a high-dimensional feature space, structural approaches represent inventions as networks comprised of *links* between more fundamental features (network nodes). In other words, structural approaches represent inventions as patterns of relations rather than as points in feature space. Network edges usually capture the relationship of components in the space of prior knowledge and practice, such as the co-presence of chemicals within a research article or technological subclasses within a patent. Nodes may also be used to represent individual inventions connected within a larger knowledge network, e.g., by co-citation (Chen et al. 2009).

Structural measures assume that the typical search for new discoveries or inventions is sufficiently incremental that existing structures in the knowledge graph (e.g., cliques or clusters) are not typically violated. These knowledge structures constitute the baseline against which novelty is established. If the discovery process is routinely slow, then a scientist or inventor familiar with the structure of this space would assess a combination of inventive elements that violates or bridges distinct clusters as less expected and more novel than combinations forged within clusters. This resonates with the claim made by Henri Poincaré (Poincaré 1952) and many since (Carnabuci and Bruggeman 2009) that the most generative combinations are “drawn from domains that are far apart” (Poincaré 1952, 24). When components under investigation are more distant in the network, more effort may be required to imagine and coordinate their combinations (Leahy et al. 2015). Hence structural approaches to novelty measurement lie literally between configurational and combinatorial approaches: They imagine invention as a process that is typically local but occasionally combinatorial across knowledge boundaries.

3.3.1. Bridging and Structural Holes.

In many structural measures, the current structure of the knowledge network is modeled at a relatively high level. A “community” of components C_i is a collection of inventive elements i, j, \dots densely connected to each other, but loosely tied to the rest of the network. In a scientific discipline, such communities might correspond to the knowledge clusters in various subfields, while in patented technology they might correspond to broad product or process categories like “telecommunications” or “pharmaceuticals” or “manufacture of plastic goods.” There are a variety of procedures for identifying communities in complex networks (Fortunato and Hric 2016), such as maximizing the modularity,¹¹ $Q = \frac{1}{2m} \sum_{ij} (A_{ij} - \frac{k_i k_j}{2m}) \delta(C_i, C_j)$, or minimizing the description length (Rosvall and Bergstrom 2010). Once knowledge clusters are identified, new links can either reinforce this prior structure (if placed within existing clusters) or violate it (if placed between existing clusters, bridging them). The former mode of invention is local and presumably less novel, especially when links connect components already joined in the past. The latter bridging mode is more novel and presumably more difficult, leading to a simple bridging measure for novelty: $Nov(i, j) = 1 - \delta(C_i, C_j)$, where novelty is assigned to the combination of inventive component i and inventive component j . $Nov(i, j)$ has value 0 when C_i and C_j are the same—components i and j are in the same cluster—and value 1 when i and j are in different clusters. For an invention D containing n components with link set $links(D) = \{(i, j), (i, k), \dots\}$, we can calculate its novelty as the number of cross community links it contains, normalized by the number of components:

$$Nov(D) = \frac{1}{n} \sum_{(i,j) \in links(D)} Nov(i, j). \quad (2)$$

These bridging approaches to novelty measurement share a basic intuition with the “structural holes” approach to social network analysis (Burt 2004). Scientists or engineers who connect components in two distinct communities span a structural hole between them; they arbitrage unexploited technical opportunities that link resources in both. Structural hole measures quantify the degree to which node or edge i spans a

¹¹Here m gives the number of edges in the network; A_{ij} are elements of the adjacency matrix representing the network, with $A_{ij} = 1$ if an edge connects i and j and 0 otherwise; k_i is the degree (number of network neighbors) for node i ; and C_i is the community to which node i belongs, with the other quantities defined *mutatis mutandis*.

structural hole by calculating constraint (Burt 2004) or betweenness centrality $Nov(i) \propto \sum_{st} \frac{n_{st}^i}{g_{st}}$, where n_{st}^i is the number of shortest paths between pairs of nodes s, t that pass through node/edge i and g_{st} counts all shortest paths between s and t (Chen et al. 2009, Cheng et al. 2010). Intuitively, a node that sits on a high fraction of shortest paths is likely to sit between and exclusively connect multiple knowledge clusters.

The components used to model knowledge structure need not be given in advance (e.g., patent codes or chemical annotations). As with distance measures, researchers have built networks out of induced, data-driven components, like the mixture of LDA topics estimated to describe a particular grant proposal. For example, Nichols (2014) argued that grant proposals featuring more rarely combined LDA topics represent novel interdisciplinary contributions.

Structural measures are especially appropriate for inventive domains where the structure of prior knowledge is important and well-organized. This organization manifests in clusters of inventive components, where most inventions draw their components from within a cluster. Consider a typical scientific paper, which must be sufficiently novel to merit publication, but nevertheless draw on components traditionally combined within an established field to gain acceptance. Such “within cluster” combinations should be more probable and less novel. Connecting things that have not been well-connected before is novel and, per Burt (2004), likely to be rewarded.

Foster et al. (2015) explored such an approach to novelty measurement in biomedical chemistry. They infer ties between chemicals co-appearing in published papers and assess the rewards associated with different forms of structural novelty. They find that bridging novelty—the introduction of new links connecting distinct knowledge clusters—positively anticipates citations and scientific awards. They found a stronger association with categorical novelty, however, which involved growing the knowledge structure by introducing new chemicals.¹² There is a similar intuition behind Verhoeven et al. (2016)’s “Novelty in Scientific Origins” indicator; in this case, they score patents based on novel connections between International Patent Classification codes and Web of Science subject categories. While the knowledge clusters are expert-identified rather than data-driven, the underlying idea is the same; they also find a positive association between structural novelty and mean impact.

3.3.2. Emergent.

A generalization of this structural approach goes beyond the transgression of prior structure by individual connections. New, *emergent* knowledge clusters violate the community partitions of the past. Such novel structures compete with and sometimes displace earlier ones. Grossman et al. (2008) use an explicit network partitioning approach to identify novel clusters of emergent behavior in an Internet IP packet network as targets for investigation. Rosvall and Bergstrom (2010) use network partitioning over journal citation networks to identify the emergence of new academic disciplines. For example, they locate the birth of neuroscience as a product of neurology fused with parts of molecular and cell biology. In the sociology of organizations and networks, Padgett and Powell’s model of emergence by autocatalysis has become a key approach in the analysis of institutional novelty (Padgett and Powell 2012), from the rise of the early modern state (Padgett and Ansell 1992) to the genesis of the biotechnology sector (Powell et al. 1996). Such measures are most appropriate when prior knowledge or invention is well-organized, structures persist over time, and longitudinal data is available for researchers to detect gradual transformations in the structure of knowledge.

3.3.3. Summary

Structural approaches assume that the search for new cultural, technological, or scientific products is gradual, so that prior structure is the essential baseline and existing knowledge clusters (typically represented through network communities) are generally respected. This is similar to the incremental search model underlying configurational approaches. Structural approaches differ in imagining the *violation* of structure

¹²They also found that new ties within an *established* knowledge cluster were more strongly rewarded than new bridges between them. Note, however, that in this case bridging ties between clusters were more common than filling in gaps within clusters, and consequently less surprising and rewarded.

as both possible and more frequent. Measures based on bridging or structural holes permit the occasional recombination of components across knowledge structures. This roughly corresponds to a long-distance move in feature space through which an invention samples from two distant feature vectors. Measures based on emergence permit something even more radical: the large-scale recombination of components to create new technologies or knowledge structures. This corresponds to the restructuring of feature space to create short cuts that accelerate discovery in new directions.

If structural approaches allow the periodic violation of prior knowledge through bridging events—slow search punctuated by faster recombination—combinatorial measures assume that the precise structure of present knowledge is largely irrelevant.

3.4. Combinatorial Novelty

Structural approaches to novelty measurement are popular among scholars of scientific knowledge (e.g., (Chen et al. 2009, Rosvall and Bergstrom 2010, Foster et al. 2015)), reflecting the importance of disciplinary knowledge structures. Scholars of technological innovation, however, favor combinatorial measures, reflecting a tradition that views invention as a recombination process (Fleming and Sorenson 2001, Arthur 2009, Verhoeven et al. 2016). Combinatorial novelty measures highlight components’ past history of recombination and the rules by which they are combined.

Combinatorial measures assume that search involves rapid recombination of inventive components. Inventors draw from the space of known components and try to assemble a valid combination. Combinations might involve chemical compounds that interact to create a robust disease therapy (Rzhetsky et al. 2015) or hardware components that interoperate to produce a useful new computer (Fleming 2001). Components are joined together as a function of their “combinability”: the likelihood that they will successfully combine with other components. Combinability is a complex function of component affordances, popularity, etc. Different combinatorial measures highlight distinct facets. New combinations between infrequent or unlikely combiners are judged unexpected and novel.

3.4.1. Peculiarity

The simplest combinatorial approach approximates the probability of a particular invention as proportional to the average combinability of its components—whether they have many distinct opportunities for combination. Fleming and Sorenson (2004) begin with the unweighted degree k_i of a component, or the number of *different and distinct* components that have been linked with it in the past. That quantity is divided by the total number of inventions in which the component has appeared, p_i , to measure the combinability of the component (e.g., a component that is connected to few distinct components, but used in many patents, is inert and *not* very combinable). The novelty of a new invention is the inverse of the average combinability of its components: $Nov(X) = \frac{n_X}{\sum_i (k_i/p_i)}$, where n_X counts the number of components in X . Inventions with high average combinability are less novel than those that break expectations by assembling a number of hard-to-combine, finicky components. We call this measure “peculiarity” because high values are achieved by assembling peculiar components: those that do not play well with others. Extensions have incorporated the date of origin for components, allowing analysts to more precisely account for recombination opportunities over shorter and longer histories (Evans 2010).

3.4.2. Atypicality

Another approach (Uzzi et al. 2013, Lee et al. 2015, Fontana et al. 2020) assesses the probability of combining *pairs* of components. If the (weighted) degree of each node w_i is treated as its combination potential and edges are randomly wired until each node’s combination potential is exhausted, then the probability that two randomly drawn components are connected is proportional to the configuration model estimate of the edge probability from random graph theory (Foster et al. 2007, Newman 2003): $Pr(i, j, t) = \frac{w_i(t)w_j(t)}{2E}$, where i and j label two nodes, w_i and w_j are their network degrees (combinability) at time t , and E represents that total number of edges in the (undirected) network. This approach is used by Uzzi et al. (2013) to model the assumption that elements frequently combined in the past are likely to recombine in the future: they are typical combinations. Uzzi et al. (2013) used journals cited in the reference list

of scientific papers as the elements of combination. They simulate a version of this random combinatorial process by fixing the network in- and out-degree of each article and randomly rewiring the citation network to establish the baseline probabilities of journal co-appearance, against which empirical co-appearances are seen as variation (quantified by a z score, with positive z scores corresponding to typical pairings and negative z scores to atypical ones). Uzzi et al. (2013) then use the 10th and 50th percentiles (“tail novelty” and “median conventionality”) of a paper’s z score distribution to characterize its blend of atypical and typical combinations of prior work, respectively. (Lee et al. 2015) introduced a related approach that substantially relaxed the computational burden of atypicality measures by using the configuration model approximation directly, rather than rewiring. In Section 4, we implement an extension of this approach to undirected networks, combining Uzzi et al. (2013)’s basic framework with Lee et al. (2015)’s computational speedup.

Uzzi’s atypicality measure can also be reformulated for computational efficiency using a vector-space calculation that interoperates well with contemporary, continuous embedding spaces (Devlin et al. 2018, Mikolov et al. 2013b). Levy and Goldberg analytically proved that pointwise mutual information (PMI), a revised z -score, equals the distance between vectorized items embedded in latent spaces as calculated by their inner product (Levy and Goldberg 2014). In this way, scientific or technological embedding spaces trained on words in technical documents or journals in reference lists learn scientific conventions and unexplored research opportunities, which can be used to facilitate exploration of scientific frontiers (Tshitoyan et al. 2019). Lin et al. (2021) demonstrate that z -score for journal pairs correlates strongly (0.74, $p < .001$) with the inner product between journal embeddings. Although this PMI reformulation is an *ex ante* measurement like the others we review, it also enables us to analyze the evolution of perceived novelty as PMI changes with a temporal sequence of embeddings (Garg et al. 2018, Kozlowski et al. 2019).

3.4.3. Obscurity

Rzhetsky et al. (2015) validated the intuition behind the combinatorial discovery model by demonstrating that most links in a network of biomedical knowledge connect drugs, proteins, and other chemicals with relatively high combinability, as measured by weighted degree. They interpret this as connecting relatively “popular” chemicals to each other. They add nuance to a basic preferential attachment picture, showing that biomedical scientists have a preference for connecting a highly combinable, popular chemical with a less popular, somewhat obscure one, in effect “building out” islands of knowledge around popular anchor chemicals—even though this strategy substantially slows the speed at which scientists collectively uncover the network of possible chemical relationships (Rzhetsky et al. 2015). In Section 4, we implement a simple version of this popularity measure; the “obscure” novelty (obscurity) of an invention D is the log of the reciprocal of the product of its weighted component degrees w_i , divided by the number of components n : $Nov(D) = -\frac{1}{n} \log(\prod_i w_i)$, where we have converted this to a surprisal measure by taking the negative log, and the product of component degrees is a proxy for the probability of the combination. Patents that score high on this measure combine many obscure, low-degree components.

In Figure 1, we illustrate how this combinatorial picture could be generalized to account for the rich combinatorial complexity of networks. For example, if we characterize a component’s combinability not by its degree, but by the combinability of its neighbors (measured through an eigenvector-based PageRank (Langville and Meyer 2006) or Bonacich Centrality measure (Bonacich 1987, Bonacich and Lloyd 2001)), then this characterization could alter which products are assessed as more and less novel. For example, components that combine with other highly combinatorial components might be transitively assumed to have higher combination potential than those that do not.

3.4.4. Multipartite Measures

Artists, inventors and scientists use heuristics beyond prior combination to assess combinability. For example, Shi et al. (2015) show that when biomedical scientists combine chemicals, diseases, methods and team members, their search is sensitive to the different probabilities $P^\alpha(i, j)$ of connecting two entities i and j through an intermediary of type $\alpha = \{\text{chemicals, diseases, methods, authors}\}$. For example, the probability of combining two previously disconnected diseases in a paper is more sensitive to connecting paths through shared methods (e.g., X-ray tomography) or chemical substances than through other diseases (Shi et al. 2015), suggesting that scientists are more attached to their methods than topics. It further

implies that when methods *are* combined because they enable better understanding of the same disease, this represents a more unexpected inventive process, improbable and novel relative to the normal pathway of discovery. To generalize, if artists, engineers, or scientists engage in any common search process, and a new discovery is viewed as unlikely to be produced by that process, the discovery will be deemed novel as a decreasing function of that likelihood. As before, this can be quantified using the surprisal $-\log P(D|K, S)$ of a particular discovery D conditioned on prior knowledge K and discovery process S .

3.4.5. Summary

Combinatorial approaches to novelty measurement assume that new discoveries and inventions are shaped more by the nature and pace of the search process than by the structure of prior knowledge and art. They also assume that search is “fast” relative to the structure of prior knowledge. Components are less characterized by their detailed positioning in some vast knowledge network, than by their combinatorial propensities.

3.5. Conceptual Summary through a Simple Model

We summarize the conceptual foundations of these novelty measures by considering a simple random walk process over the space of components and the inventions they compose. An appropriate random walk process can be defined for configurational, structural and combinatorial measures and their underlying representations (feature spaces, knowledge networks, etc.).

Suppose that the typical scientist, inventor, or designer discovers new products by drifting randomly through the network of known inventive components, testing and rejecting possible links. The more unlikely it is that a scientist exploring at random will reach from one inventive component i to another j , the more difficult it is to uncover a valid combination between them (Chen et al. 2009). A random walk will tend to stay within a local neighborhood of densely connected ties, rather than venturing to loosely connected domains, allowing it to capture structural features (see Rosvall and Bergstrom 2010, ’s approach to community detection).

A random walk interpolates between local, configuration and structure-based search, on the one hand, and non-local, combinatorial search, on the other. To keep track of how many times elements have been combined in previous inventions, we encode the network of inventive combination in a weighted adjacency matrix. At time t , the novelty of a combination (i, j) is a decreasing function of the probability to diffuse from i to j at t . Formally, if $P_{ij} = \frac{W_{ij}}{\sum_{r=1}^m W_{ir}}$, then the probability for a random walker originating at component i to be at component j after n steps is $Pr(i, j|n) = (\mathbf{P}^n)_{ij}$. Now consider the continuous-time version in which $X(t)$ provides the position of the random walker at time t :

$$Pr(X(t) = j|X(0) = i) = \left[\sum_{h=1}^{\infty} \frac{(\lambda t)^h e^{-\lambda t}}{h!} \mathbf{P}^h \right]_{ij} = [e^{-Qt}]_{ij} \quad (3)$$

where λ is the diffusion rate and λt fixes the average number of steps h taken in time t (Montroll and Scher 1973). We can estimate the most likely diffusion parameter λ , given time-stamped data, via maximum likelihood procedures (Cokol et al. 2005:200; Gilks, Richardson, and Spiegelhalter 1996) to parameterize the scientific discovery process. Random walks with a large λ describe a rapid exploration process, largely insensitive to distance and structure, and as $\lambda \rightarrow \infty$, the probability of landing on a specific element becomes proportional to its degree¹³. Assuming that components are chosen by periodically stopping this high speed random walk, then the probability of reaching between two components will be proportional to the product of their degrees, which mimics combinatorial novelty measures like those proposed in (Fleming and Sorenson 2004) and (Uzzi et al. 2013). By contrast, random walks with a small λ (e.g., $\lambda = 0.01$) will explore slowly,

¹³Technically, the stationary distribution of the random walk process will be proportional to the degree of each node, such that the probability for a walk begun from any node i to end at another node j will be proportional to the degree of j , so long as i and j are in the same connected component.

making local connections highly probable, just as in bridging or emergence models. We can then define the novelty of combinations as some decreasing function of their probability, as in Equation 1¹⁴.

We note that random walk measures are computationally expensive relative to existing approaches. We propose the model as a device for illustrating the relationship between novelty measures and showing how they all rely implicitly or explicitly on an underlying model of the discovery process. Insofar as our random walk model, under different λ s, corresponds to our distinct novelty measures, it will confirm the accuracy of our characterization of those measures as simulations of search processes with varying speed. Insofar as our measures reveal variations in the inventive process across technological sectors, it will justify the importance of adapting measurement to sector-specific search processes and knowledge structures in order to capture the perceived novelty of any new invention. We turn to such variation in the next section, testing our framework on the analysis of novelty in U.S. Patents.

4. Testing the Framework: Novelty Measures in U.S. Patents

Here we use inventions disclosed in U.S. patents to demonstrate the value of our framework for novelty measurement. We draw on data from approximately 2.5 million U.S. utility patents, accessed through Patentsview¹⁵. We represent prior knowledge using the 2,259,119 patents granted between 1976 and 2000; and assess the novelty of 157,595 patents granted in 2000. To check the robustness of our findings, we repeat our analyses on 90,421 patents granted in 1990, representing prior knowledge with the 958,817 patents granted between 1976 and 1990. Findings are broadly similar across these two cases.

We use codes from the United States Patent Classification (USPC) system as basic invention components. The USPC is a two-layer classification system. The top layer consists of terms called classes. According to USPTO, a class generally delineates one technology from another. Every patent is assigned a main class. Each class also contains subcomponents called subclasses. Subclasses delineate processes, structural features, and functional features of the subject matter encompassed within the scope of a class. In total, there are 496 class codes and 158,073 subclass codes.

We implement 5 novelty measures drawn from the three categories discussed above, and calculate the novelty scores for each patent in our population of interest. First, we show that technology areas vary in the way novelty measures are distributed. This result follows from our claim that different novelty measures rest on different ‘conceptions’ of the inventive process. Second, we document the relationship between different novelty measures and the mean, tail, and variance of impact in each technology (i.e., class) as well as higher-order groupings of technologies associated with USPTO “technology centers.” We find that technologies and technology centers vary widely in which novelty measures predict impact. More importantly, we find that technology fields closer in inventive practice have much more similar relationships between novelty and impact than those further apart. Given well-established associations between novelty and impact, e.g., (Fleming 2001, Uzzi et al. 2013, Foster et al. 2015, Lee et al. 2015), this suggests that similar inventive processes are best captured by similar novelty measures. Finally, we show that our simple random walk model tracks associated novelty measures as walk speed is varied, justifying our characterization of inventive domains and associated models as fast or slow.

4.1. Novelty Measures Used

We implement five measures from the literature as illustrative examples of the configurational, structural, and combinatorial approaches to novelty measurement. Brief definitions of the measures are given below; see Section 3 for detailed exposition and discussion.

¹⁴ λ can be estimated even more precisely; we need not assume that the adjacency matrix is static. For example, we could divide the time interval T into intervals Δt_s during which the weighted adjacency matrix is fixed as $\mathbf{W}(s)$, but at the end of each time interval, a new adjacency matrix is used, taking into account new combinations discovered in the intervening interval. The probability that a continuous-time random walk starting at node i will be at node j at time $t = T$ is then given by $Pr(X(T) = j | X(0) = i) = [\prod e^{-\mathbf{Q}(s)\Delta t_s}]_{ij}$ where $\sum_i \Delta t_i = T$.

¹⁵<https://www.patentsview.org/>

Configurational novelty. Following the literature on patents (Fleming 2001), we use subclasses from the USPC system as features for specific patents. Each patent X is thus represented by a feature vector $f^X = (f_1^X, \dots, f_n^X)$ of subclasses, where $f_i^X = 1$ if patent X is associated with subclass i and $f_i^X = 0$ otherwise. The simple, configurational novelty of patent X is defined as the median Hamming distance from X to all patents Y granted prior to X in the same main class:

$$Nov(X) = Median\{\sum_i |f_i^X - f_i^Y|, \forall Y \text{ granted prior to } X \text{ in the same class}\}.$$

This measure is common in the biology literature (Gavrilets 1999, Foster et al. 2013a, Erwin 2020); it is closely related to novelty measures based on cosine distance (Nanda et al. 2013), and variants have been used for patent classification (Seneviratne et al. 2015).

For other novelty measures, we construct a co-occurrence network of subclasses to represent the underlying search space. Two subclasses are connected if they have ever co-appeared in a patent granted before 2000; the link is weighted by the number of patents in which they co-appear. Accordingly, a patent is represented not only by the subclasses associated with a patent, but also by pairwise links between those subclasses. A specific patent is, in essence, a clique (complete graph) across its subclasses, as in (Foster et al. 2015, Verhoeven et al. 2016).

Structural novelty. We use the Louvain method (Blondel et al. 2008) to find communities in the subclass network. A patent’s structural novelty is calculated as the number of cross-community links in the patent, normalized by its total count of subclasses (Equation 2). This measure is a continuous version of the one in (Foster et al. 2015), and a close conceptual analog of the “scientific origins” measure in (Verhoeven et al. 2016)¹⁶.

Atypicality, our first combinatorial measure, captures how “unusual” the combinations in an inventive product are. It is defined as the negative of the 10th percentile of z-scores for combinations in the inventive product, as in (Uzzi et al. 2013). Recall that atypical combinations have a negative z-score; a product with several unusual combinations will have a large negative value for the 10th percentile of its z-scores and thus large atypicality. The measure developed in (Uzzi et al. 2013) requires a few modifications to be used on the patent network. Links represent co-appearances between patent subclasses instead of citations between papers, and hence are undirected. Consequently, we can calculate the z-score of a link (k, l) directly as $\frac{w_{kl} - E[w_{kl}]}{\sigma(w_{kl})}$ (where w_{kl} is the weight of link (k, l) , $E[w_{kl}]$ and $\sigma(w_{kl})$ are the expectation and standard deviation of w_{kl} under the configuration model¹⁷; see section 3.4.2 for detailed discussions). We then assign the novelty of a patent as the negative 10th percentile of z-scores for all links in the patent¹⁸.

Obscurity, our second combinatorial measure, is defined as $-\sum_k \log(w_k)/n_X$, where w_k is the weighted degree of subclass k associated with patent X . The normalizing factor n_X is the number of subclasses in X .

Pecularity, our third combinatorial measure, is an exact replica of the well-known measure used in (Fleming and Sorenson 2004). It is defined as $n_X / \sum_k (d_k/p_k)$, where d_k is the unweighted degree of subclass k and p_k is the number of patents already associated with subclass k .

4.2. Distribution of Measures Across Main Classes

In Figure 2, we plot the distributions of different patent novelties for technology classes drawn from 8 of the 9 USPTO “technology centers”. Each technology center examines patents across a range of related

¹⁶Our version is normalized, uses data-driven knowledge clusters, and examines connections within technological clusters rather than between technological and scientific knowledge. Our approach finds 12 substantial communities in prior knowledge from 1976-2000, which map onto broad technical domains, as well as four small and less coherent communities. We provide a brief gloss of the coherent communities here: traditional manufacturing and industry; telecommunications; electricity/electrical; industrial chemistry, heating, and refrigeration; fine metal manufacture; plastic goods; pharmaceuticals and biomedicine; computer hardware; resins, rubbers, and coatings; molecular biology, microbiology, and food; ammunition and explosives; and (amusingly) beekeeping.

¹⁷With mean-field approximation, $E[w_{kl}] \approx \frac{s_k s_l}{m}$ where s_k is the strength (i.e., weighted degree) of k and m is the total strength of the network, and $\sigma(w_{kl}) \approx \frac{s_k s_l / m}{1 - 2s_k s_l / m^2}$

¹⁸This strategy is very similar to the approach of Lee et al. (2015), who likewise use a configuration-model approximation, but do not use it to approximate z-scores as we do here.

classes; by looking at a typical class from each center, we are effectively looking across the landscape of inventive domains, as seen from the USPTO. We exclude the center focused on design patents, as their use of class annotations differs from utility patents¹⁹. We also exclude design patents and design classes, in the analyses below, for the same reason.

Across technology classes and novelty measures, most distributions take on a rough inverse-U shape. This is consistent with the intuition that most patents are “normal,” with very few being highly novel or trivial. Marked differences are revealed between fields, however. For example, inventions in the class “502. Catalyst, solid sorbent or support therefor” exhibit great variation in structural novelty, suggesting that invention in this domain often involves the transgression of knowledge clusters. By contrast, patents from “370. Multiplex Communication” show *little* such variation, implying that invention in this domain almost entirely respects knowledge cluster boundaries, which is unsurprising as one of the clusters corresponds to telecommunications. Several classes have broad distributions of Peculiar novelty, suggesting that some patents put together several “hard-to-combine” components). Others have tightly peaked distributions, meaning that patents in these domains tend not to vary much in the combinability of their components.

To verify that these differences are not due to randomness or measurement noise, we perform a series of statistical tests. First, for each novelty measure, we perform an ANOVA between all 473 classes²⁰ to test the difference in the first moments. As a robustness check, we then perform the Kruskal test on our measures to test differences in their medians. Similarly, we perform both Levene’s test and the Brown-Forsythe test for the second moments (or variance) of the measures between classes. All results are significant at $p < .001$ (see Appendix Table B.2 for detailed results), rejecting the null hypothesis that the measures have equal mean, median, or variance across the classes.

While these novelty measures have been successfully deployed to measure various inventive processes, the different distributions shown in Figure 2 underline that they measure different aspects of invention and its relationship to prior knowledge. This result follows from our core argument; these measures have been developed for specific research questions, rely on distinct assumptions, and do not trivially generalize to different domains (even within the space of technology). Insofar as different technology areas have distinct inventive processes and differently structured prior knowledge, we would expect their patents to exhibit different ranges and distributions of novelty across various measures. That is exactly what we see here.

4.3. Novelty Measures and Impact

Novelty measures are often used to characterize and predict influential discoveries, inventions, and products. Indeed, the use of *ex ante* novelty measures to identify potential breakthroughs and blockbusters is one of the primary reasons to study novelty. Deep characterization of the association between novelty and impact contributes directly to the study and management of invention and innovation, as well as the crafting of science and technology policy. Here we explore the relationship between distinct novelty measures and patent impact, using citation as a proxy for impact. All evaluated patents are published the same year (2000), and we count the citations they have accumulated from 2000 to 2010 from other granted patents²¹. Although of course a simplification, citation is still the best widely available quantitative measure of impact. It has a demonstrated association with the diffusion of ideas (Gerow et al. 2018) and commercial value (Harhoff et al. 1999). It is also widely used in innovation studies; this renders our analysis easily comparable with the existing literature, including very recent articles like (Fontana et al. 2020).

The novelty framework described above motivates the following deduction. Because different inventive areas have distinctive invention processes, different measures will provide the best characterization of novelty in those domains. Because invention, like innovation (Cohen et al. 2019), is fundamentally multidimensional, a combination of measures will characterize novelty better than a single measure. Since novelty is associated with impact, it follows that the relationship between novelty measures and different forms of impact should also vary across inventive areas, whether single or multiple measures are used.

¹⁹We thank the editor for this critical point, noting also that including them does not change the pattern of results

²⁰Excluding design classes.

²¹Our data end in 2010. In our robustness test, we also examine patents published in 1990 and count the citations they have accumulated from 1990 to 2010. These tests confirm the findings described here.

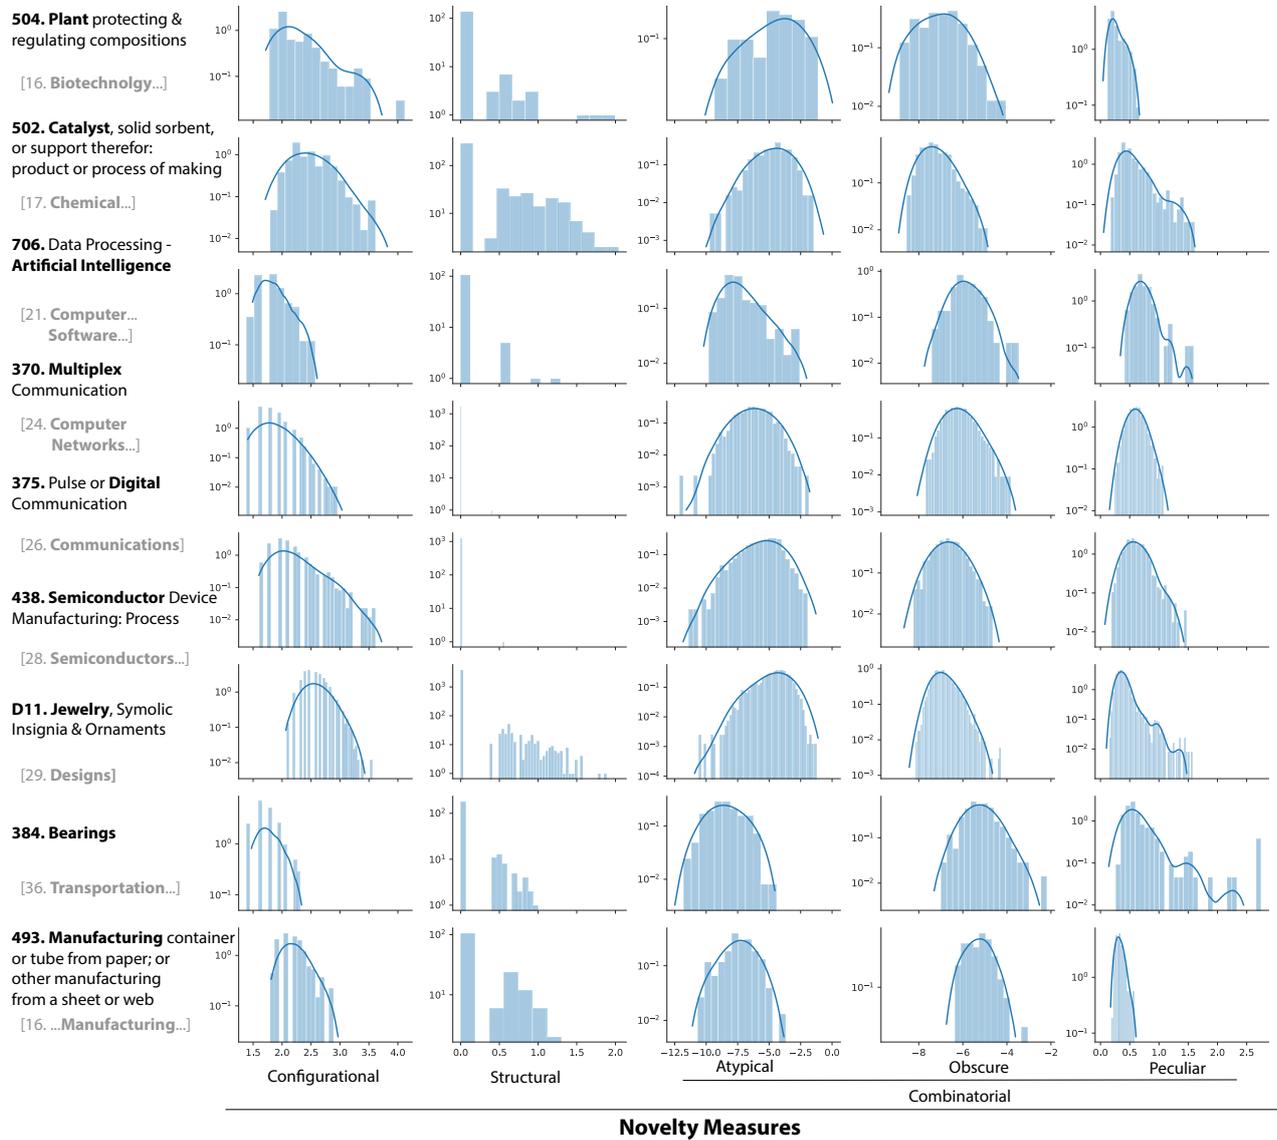


Figure 2: Distinct technology classes (in black) from different USPTO “technology centers” (in gray) reveal marked differences in the distribution of novelty measures. Solid lines are kernel density estimates with Gaussian kernels.

To test this prediction of our framework, we build three classes of regression model that predict mean impact, hit impact and the variance of impact, using the five novelty measures implemented above as predictors²². All three variants of the novelty-impact relationship are prominently attested in the literature, e.g., (Foster et al. 2015, Uzzi et al. 2013, Fleming 2001) for mean, hit, and variance of impact respectively. To capture potential nonlinear relationships between novelty and impact, as widely reported in the literature (Uzzi et al. 2013, Foster et al. 2015), we also include the square of each term as predictors.

The first class of models regress the logged count of citations received by a patent on all five novelty measures and their squares; a distinct model is estimated for each patent class. The second model uses logistic regression to predict whether a patent is among the top 10% most cited patents granted in 2000 (resp. 1990) within each class; this model uses the five novelty measures and their squares as predictors, and is also estimated for each patent class. The third takes a multi-level modeling approach to regress the variance of citations on all five novelty measures. We assume that the logged count of citations for each patent x has a Normal distribution: $\log(C_x) \sim N(\mu_x, \sigma_x^2)$, and both mean and variance of this distribution are functions of the novelty measures: $\mu_x = a_0 + \sum_{i=1}^5 a_i Nov_i(x) + \sum_{i=1}^5 a'_i Nov_i(x)^2$ and $\sigma_x^2 = \exp(b_0 + \sum_{i=1}^5 b_i Nov_i(x) + \sum_{i=1}^5 b'_i Nov_i(x)^2)$ where $Nov_i(x)$ is the i -th novelty measure of patent x . Our goal is to find parameters a_i, a'_i, b_i, b'_i that maximize the likelihood of observing the log-citation counts in our data given their novelty scores.

For each model, the estimated coefficients for the novelty measures reveal the differential impact of each implicit novelty-generating process on the reception of patents in that technology class. As a result, we call a particular configuration of coefficients the *impact profile* of the novelty measures for that class.

We fit our models to each class of patents separately. In these analyses, we focus on how novelty characteristics of patents *within a particular inventive area* predict their impact; we therefore compare patents within the same class and the same year. We do not control for other patent characteristics (e.g., number of inventors, assignee, etc.). While these no doubt have some relationship with impact, they also affect the novelty characteristics of the patent (Lee et al. 2015, Wu et al. 2019), and may moderate the relationship between novelty and impact (Lee et al. 2015). Such questions are, of course, interesting, but orthogonal to our purpose: To ask whether the core relationship between novelty and impact varies across measures and inventive fields.

We can use these models to test a core prediction of our novelty framework. When two classes frequently cite each other, they are likely to have similar inventive practices (Yan and Luo 2017). If the inventive practices in two classes are similar, then the distinct novelty-generating processes captured by each measure should have a similar net effect on impact. This implies that the impact profiles of the two classes should be similar as well. By contrast, when two classes *rarely* cite each other, their inventive processes are likely to be quite different—implying different impact profiles.

We show this analysis in Figure 3. The horizontal axis indicates the cosine similarity between two impact profiles, with 1.00 indicating identical profiles and -1.00 indicating “opposite” profiles. The vertical axis indicates the cross-citation strength between the two classes (i.e., the sum of log-citation counts between the two classes) for which the corresponding impact profiles are calculated. For clarity of presentation, we omit individual data points here and show fitted trend lines; we also remove the intercept terms of the fitted lines to accommodate them on the same plot. (Plots including the raw data points and individual fitted lines are included in the Appendix.) The shaded area around each line indicates the 95% confidence intervals.

While the change of y -values in Fig. 3 does not seem dramatic, the trends are statistically significant: We find a significant and positive association between impact profile similarity and cross-citation strength. This suggests that fields sharing technologies and knowledge also reward and value the same novelty-generating processes—just as our framework predicts. We note that the impact profile similarity manifests the strongest association with the variance of impact (correlation $\rho = 0.32, p < .001$) and weaker associations with mean impact ($\rho = 0.08, p < .001$) and hit impact ($\rho = 0.04, p < .001$).

Our second analysis examines the relationship between novelty and impact within each of the 8 USPTO technology centers²³. We first repeat the three families of regressions described above for each class, but

²²We also repeat these analyses using each novelty measure as a single predictor; see below.

²³A technology center examines patents across multiple classes as specified by the USPTO; broadly speaking, centers examine

use only one measure at a time as a predictor (including the squared term as above to model non-linearity reported in the literature). We then collect the R^2 and pseudo- R^2 for each measure-class combination. For all classes within a technology center, we compute the fraction of the classes for which a given measure is the best predictor; this information is encoded in Figure 4 as a heat map. Very dark colors mean that a large fraction of classes in a particular Technology Center have that measure as the best predictor (e.g., configurational novelty is the best predictor of mean success for the vast majority of classes in Tech Center 2100 (Computer Architecture Software and Information Security)); very light colors mean that a small fraction of classes have that measure as the best predictor (e.g., structural novelty is rarely the best predictor of hit success for Tech Center 2100).

In each cell, we also show the average R^2 for that particular measure; the average is taken across all classes in the appropriate technology center. While these numbers are typically quite small, we emphasize that the average includes cases where a particular measure is the *best* predictor as well as cases where it is the *worst*. For many measure-class combinations the R^2 can exceed 0.1 or 0.2. This underlines the importance of choosing novelty measures wisely; it also suggests that some inventive areas have a weaker link between novelty and impact, across a range of measures, a point that we pursue below.

We also see that technology centers differ substantially in how different facets of novelty predict impact. For example, configurational novelty is important for predicting mean and hit success across most centers; it is especially important for Tech Centers related to Computing and Communications (21, 24, and 26). Even in the same center, impact profiles differ depending on which form of impact is predicted. For the Biotechnology and Organic center (16), configurational and structural novelty are both often good predictors of hit success (consistent with Foster et al. (2015)’s findings about bridging knowledge clusters), while obscurity is the best predictor of variance success, consistent with the finding that highly cited and prize-winning work tends to combine less popular entities in this domain (Rzhetsky et al. 2015).

This variation—across centers and even across kinds of impact—provides further support for our assertion that, because different inventive domains possess differently structured prior knowledge and distinct inventive processes, they should vary in which novelty measures best capture the perceived novelty of inventions in that domain, and hence predict well-known correlates of novelty like impact.

4.4. Conceptual Summary with Random Walks

Having provided an empirical illustration of our major conceptual points, we next evaluate our conceptual summary of the organization of novelty measures with the random walk model (Section 3.5) over the network of subclasses. As before, the weighted network of subclasses is constructed from patents granted prior to 2000. We assess the novelty of patents granted in 2000 and assume that $t = 1$ in the random walk model (Equation 3). Recall that λ is the “speed” of the random walk and parameterizes the rate of the discovery process. Larger λ means more rapid walking and hence faster speeds of search; as λ gets larger, the discovery process should shift from more configurational to more combinatorial. The random walk novelty of patent X is defined as $-\sum_{i,j} \log(Pr(j|i) + Pr(i|j))/E_X$, the average random walk novelty over all pairs i, j in the patent²⁴.

For each patent, we calculate its random walk novelty at 30 random walk speeds from 0.01 to 0.1. We then calculate the mutual information between the random walk novelty at each speed and every other novelty measure across all patents. Mutual information is an information-theoretical measure of the mutual dependence between the two distributions: It is zero if the two distributions are independent and non-negative in general. A high mutual information between two novelty measures (e.g., random walk novelty and obscurity novelty) means that they have similar distributions across patents. Compared to the correlation coefficient, it is not limited to linear dependence but it is harder to compute. Formally, the mutual information between two variables X and Y is defined as the Kullback - Leibler divergence between the joint distribution $P(X, Y)$ of the two variables and the product $P(X)P(Y)$ of the marginal distributions. In our case, X would be one of

patents that are in the same areas of technology, e.g., “Biotechnology and Organic” (Tech Center 1600) or “Semiconductors, Electrical and Optical Systems and Components” (Tech Center 2800).

²⁴ E_X is the total number of pairs of nodes in patent X

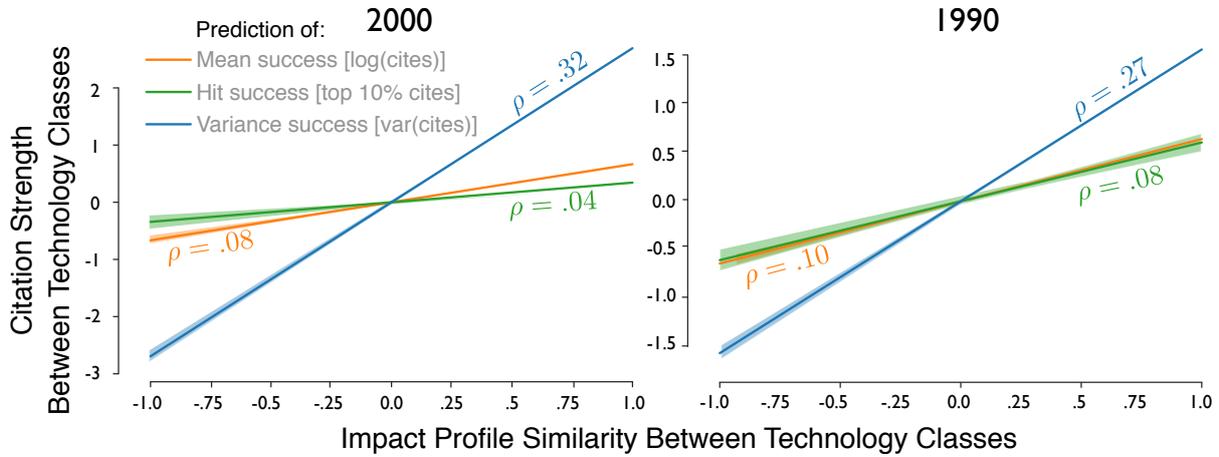


Figure 3: Technology classes that cite each other have similar impact profiles. For each measure of impact, we regress the cross-citation strength between a pair of classes (Log of the union of citations from class i to j and from class j to i) on the cosine similarity between the impact profiles of the two classes (i.e., a simple linear regression). For each class, we compute three impact profiles (mean citation, hit citation, or variance of citation) and hence estimate three regression models (denoted by orange, green, and blue colors respectively) for patents from year 2000 and also 1990. Results suggest that similar technology areas (indicated by extensive cross-citation) share novelty-generating processes and show a similar relationship between novelty processes and the reception of inventions. For clarity of presentation, we omit individual data points but show fitted trend lines with 95% confidence intervals (the data points are shown in Fig. A.1). Beside linear relationships, we show the Pearson correlation coefficients ρ linking impact profile similarity and citation strength for pairs of technology classes. All results are significant at the .001 level in both years.

USPTO Technology Centers	Average R^2 for Novelty to Predict																			
	Mean Success log (citations)					Hit Success top 10% citations					Variance Success var (citations)									
	0.0	.15	.30	.45	.60	.75	.90	0.0	.05	.10	.15	.20	.25	.30	.35	0.0	.10	.20	.30	.40
16. Biotechnology & Organic	0.039	0.052	0.039	0.031	0.048	0.035	0.038	0.023	0.022	0.027	0.025	0.036	0.029	0.024	0.036					
17. Chemical & Materials Engineering	0.044	0.034	0.034	0.035	0.038	0.022	0.016	0.018	0.015	0.022	0.020	0.037	0.019	0.019	0.022					
21. Computer Architecture, Software & Information Security	0.043	0.010	0.027	0.027	0.030	0.034	0.021	0.026	0.027	0.029	0.023	0.019	0.014	0.016	0.034					
24. Computer Networks, Multiplex, Cable & Cryptography/Sec.	0.042	0.006	0.020	0.025	0.022	0.025	0.011	0.011	0.017	0.017	0.015	0.025	0.008	0.010	0.009					
26. Communications	0.035	0.005	0.021	0.022	0.044	0.018	0.004	0.019	0.010	0.023	0.010	0.010	0.009	0.007	0.011					
28. Semiconductors, Elec. & Optical Systems & Components.	0.042	0.024	0.033	0.032	0.025	0.020	0.017	0.024	0.020	0.020	0.027	0.061	0.026	0.022	0.028					
36. Transportation, E.Commerce, Constr., Agr., Licensing & Review	0.045	0.037	0.039	0.059	0.041	0.020	0.017	0.021	0.023	0.025	0.035	0.096	0.034	0.051	0.064					
37. Mechanical Engineering, Manufacturing & Products	0.046	0.040	0.041	0.052	0.058	0.017	0.027	0.031	0.028	0.045	0.038	0.084	0.040	0.040	0.044					
	configurational	structural	atypical	obscure	peculiar	configurational	structural	atypical	obscure	peculiar	configurational	structural	atypical	obscure	peculiar					
		combinatorial					combinatorial					combinatorial								

Figure 4: Average estimated rank of novelty measurements in their prediction of well-studied novelty correlates—mean, variance and hit success—for each of the 8 USPTO specified “technology centers.”

the novelty measures and Y the random walk novelty measure at a certain speed. However, we do not know the distributions of the novelty measures but only observe samples from the distributions (i.e., the novelty scores of the patents); therefore, we take a non-parametric approach that locally estimates the probability density at each sample point, similar to kernel density estimation. For implementation details, please refer to the Non-parametric Entropy Estimation Toolbox²⁵, which is used to calculate the mutual information scores in this analysis. In Figure 5, we show how mutual information between the random walk novelty and every other novelty measure varies as the walk speed increases. We also perform a linear regression of the mutual information on the random walk speed for each novelty measure. The regression analysis tests the null hypothesis that the mutual information between a given novelty measure and the random walk novelty does not change with the walk speed. A significant coefficient of the walk speed then provides evidence to reject the null hypothesis. For configurational and structural novelty, mutual information decreases as random walk speed increases, with configurational novelty dropping the fastest ($\beta = -0.0043$, $p < .001$) and structural novelty dropping gradually ($\beta = -0.0005$, $p = 0.025$). For combinatorial novelties, mutual information increases as random walk speed increases ($\beta_{peculiarity} = 0.0049$, $p < .001$; $\beta_{atypicality} = 0.0065$, $p < .001$; $\beta_{obscurity} = 0.0120$, $p < .001$). This pattern further substantiates our conceptual framework, and our underlying claim that configurational novelty embodies a slow, local search process while combinatorial novelty assumes a fast, non-local one.

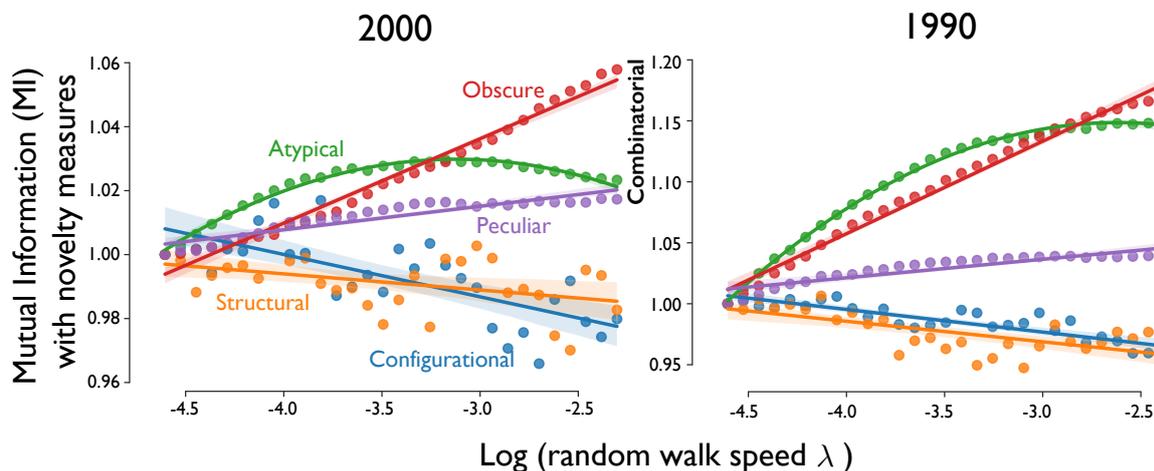


Figure 5: Mutual information (MI) between simulations of invention with different random walk speeds (λ) and five common measures of novelty; MI normalized for $\lambda = 0.01$ to fit all data points into the same plot. Regression lines (and 2nd order polynomial for Atypicality) are fitted to the data points with 95% confidence intervals to aid visualization of the significant association between mutual information and random walk speed.

4.5. Summary and Methodological Implications

Our empirical analysis demonstrates the basic validity of our conceptualization of novelty: Novelty is a function of surprise under differing models of invention, and surprise is conditioned on the prevailing inventive process S and the organization of existing knowledge K . Our findings here validate several key predictions of this framework. First, we find that patents from different technology areas exhibit different ranges and distributions of novelty across novelty measures. This is just what we would expect if these areas have distinct inventive processes and differently structured prior knowledge, which different measures capture more or less adequately. Second, we find that when patent classes frequently cite each other, they have similar impact profiles, i.e., relationships between novelty measures and impact. Again, this follows

²⁵<https://github.com/gregversteeg/NPEET>

directly from our framework: They cite each other because they have affinity in inventive processes, practices, and underlying knowledge; under our framework, we would expect similar novelty measures to do a better job of capturing that novelty. Third, we find pervasive variation across technology centers in which novelty measures best predict impact, underlying our assertion that differently structured knowledge and distinctive inventive processes imply different models of discovery and hence different novelty measures.

We emphasize that this validation of our conceptual framework has direct, and striking, implications for research practice. First, it sinks the dream of a single, “silver-bullet” measure of novelty. This means that researchers should put considerable care and effort into selecting and justifying their novelty measure as ecologically valid. In addition to predictive analyses of the kind we perform here, researchers should consider experimental methods to validate measures, or draw more heavily on ethnographic and other qualitative evidence from their chosen inventive domain (Azoulay et al. (2017) provide a nice example of using qualitative evidence to justify assumptions of quantitative methods). Second, this analysis (and our broader framework) suggests that scholars of innovation should re-dedicate themselves to developing more nuanced mathematical and computational representations of knowledge K and the invention process S , from which they can build better models of discovery and hence better measures of novelty. We take up this challenge in the next section. Here, too, the space for valuable qualitative research is huge, as such work can capture the texture of K and S in specific domains.

5. New Directions in Novelty Measurement

Having reviewed the existing novelty measures and empirically illustrated our framework and key conceptual claims with patent data, we now turn to promising new directions in novelty measurement: better characterization of complex and evolving search spaces (Section 5.1) and variation in subjective surprise (Section 5.2).

5.1. Complex and Evolving Search Spaces

Table 1 and Figure 1 highlight considerable internal variation within configurational, structural and combinatorial approaches to novelty measurement. This variation tracks the complexity of the search space involved. Early approaches to novelty measurement assumed a simple search space built from individual words, citations, pixels, or sub-technologies (e.g., patent subclasses) and their co-appearance in cultural products. In such approaches, any components at the same configurational, structural or combinatorial distance are equally likely to connect. More recent approaches allow researchers to build more distributional (Murdock et al. 2015) or structural information (Gross 2014, Shi et al. 2015) into novelty measures.

One way to add structural information is by including more *types* of components. This has changed our understanding of inventive processes and improved novelty prediction. Consider work by Pontikes (2014), which analyzes the relationship between patent novelty and the venture financing that start-ups obtain. Better financing is driven not by novel combinations of technical components within firm patents, but rather by novel combinations of technical components with new market categories. Shi et al. (2015) analyzed 6 million recent MEDLINE abstracts to understand how biomedical scientists weave new links in the network of knowledge, taking into account not only the identity of each node, but also its type (e.g., diseases, method, chemical). The structure of the inventive space changes considerably once many component types are included, becoming extremely dense and making simple network measures uninformative. Moreover, connections between any one type of component (e.g., between two diseases) were much more likely to be mediated by a distinct type, such as a new method or chemical signature that brought them together. These examples demonstrate how analyses that include more component types, and allow for their complex combination, can better predict the growth of new technology and better approximate the process through which scientists and inventors think.

Striking discoveries are often perceived as glaringly obvious by those who did not make them. This is because inventors use features that are *known* but not cognitively available to most practitioners (Tversky and Kahneman 1974). Such discoveries are seen as “reframing” a field. The logic of the breakthrough invention makes suddenly salient a new feature or category of features, restructuring the relationship between existing components (Kuhn 1962). Howard Margolis advances a similar proposal in *Paradigms and*

Barriers (Margolis 1993), arguing that “habits of mind” are the primary barriers to revolutionary scientific discovery. *Creative* invention, discovery and production often involve recombination through new patterns of association. We believe that identifying such reframing moments—the introduction of novel principles of discovery—is an especially promising direction for scholars of innovation, and one that requires complex, nuanced representations of the underlying search space as it is rewired by such principles.

The structure of complex search spaces can also be transformed by changes in its constituent components. Several prominent theories of invention and innovation lean heavily on the idea that over time, groups of components used in concert come to be seen as a single unit (Arthur 2009, Latour 1987). Brian Arthur’s landmark analysis of technology details a process by which inventors combine and then miniaturize objects into stable complex components, which become the units of selection and recombination for future inventors (Arthur 2009). Through frequent reuse, these new complex components are eventually named. Consider the “motherboard”, a circuit board that connects the core electronic components of a personal computer²⁶. Before the physical integration of the motherboard, the backplane (and before that, a system of cables) connected the same subcomponents, but with less specialized or forced coupling between them. Bruno Latour describes a related process termed “black boxing” (Latour 1987), which facilitates the reuse of complex components in new contexts. Black boxing typically involves a differentiation of the internal and external layers of the emerging complex component, with most sub-components exclusively linked *inside* the complex component and only a few sub-components forming a robust *interface* that connects inside with outside (Simon 1969). The input/output (I/O) interface for each complex component among a field of such components may stabilize into a widely used protocol (e.g., USB, URL, syringe injection). Interoperability with such protocols increases the likelihood that a component will be reused and become central to the inventive field. Hence new complex components typically use I/O that obeys one or a few common protocols for interacting with their “environment” (Galloway and Thacker 2007), unless they are part of an ambitious bid to transform that environment (e.g., Apple Computer). When complex components are black-boxed and manufactured as products, novelty measurement comes full circle. Complex components, now seen as irreducible from their function, come to be perceived and treated as simple, categorical novelties—fundamentally new. These emergent components then enter the vocabulary of features used in novelty measurement. For example, the US and European patent offices annually issue additional and reorganized patent subclasses to correspond to frequently used simple or complex components.

Understanding the complex components that represent the frontier of knowledge in an inventive domain allows us to characterize the “adjacent possible” (Kauffman 2000), an evolutionary concept widely applied to innovation in science and technology (Johnson 2011, Tria et al. 2014, Wagner and Rosen 2014). The adjacent possible comprises all things (e.g., molecules, ideas, technologies, or art) that lie a single “step” beyond whatever exists currently. Whenever a novel invention, discovery, or cultural work emerges, the adjacent possible expands to embrace a fresh set of proximate possibilities²⁷. We propose that dimension reduction methods can be used to “discover” complex components visible to those in the inventive and evaluative fields but not yet codified in component ontologies (Bengio et al. 2013). For example, matrix factorization approaches like principle components analysis (PCA), neural network autoencoders (Goodfellow et al. 2016), or stochastic block models (Decelle et al. 2011) could be used on an adjacency matrix of inventive components within sequential time windows to reveal the evolution of dominant, low-rank components in the creative space. This dynamic approach to feature creation could be used to pre-process data for novelty measurement. A more ambitious model could directly incorporate dimension-reduction into its simulation of novelty creation.

5.2. Measuring Subjective Surprise

The novelty measures reviewed in Section 3 connect novelty to improbability and surprise, but they assume a shared subjective assessment of an invention’s probability. This is clearly problematic. Variation in

²⁶Here we draw on the excellent summary in <https://en.wikipedia.org/wiki/Motherboard>.

²⁷Consider how the capacity to “plug” a complex component into an existing configuration of other components as a single step radically transforms the nature of the adjacent possible.

the experiences and beliefs of the agent simulating discovery and assessing novelty have clear consequences for her assessment of novelty (Guinan et al. 2013, Lamont 2009). To capture this variation, we begin with a more general, Bayesian approach to surprisal measurement, in which prior beliefs are explicitly modeled. Bayesian Surprise incorporates expectations regarding an outcome—the subject’s “model”—and then, drawing on Bayes Theorem, estimates the degree to which the subjective model must be updated to account for new experiences.

Quoting its inventors in Itti and Baldi (2005), “the background information of an observer is captured by his/her/its prior probability distribution $\{P(M)\}_{M \in \mathcal{M}}$ over the hypotheses or models M in a model space \mathcal{M} . Given this prior distribution of beliefs, the fundamental effect of a new data observation D on the observer is to change the prior distribution $\{P(M)\}_{M \in \mathcal{M}}$ into the posterior distribution $\{P(M|D)\}_{M \in \mathcal{M}}$ via Bayes theorem, whereby $\forall M$ in \mathcal{M} ”: $P(M|D) = \frac{P(D|M)}{P(D)} P(M)$.

With information-theoretic surprisal, an observation is unsurprising if it has high probability. With Bayesian Surprise, an observation is unsurprising if it requires little updating of the prior, i.e., if the support $\frac{P(D|M)}{P(D)}$ is close to 1. With information-theoretic surprisal, an observation is surprising if it has low probability. With Bayesian Surprise, it is surprising if it requires a substantial update of the prior—if the support differs substantially from 1. We can quantify the change between prior and posterior using the KL divergence:

$$S(D, \mathcal{M}) = KL(P(M|D), P(M)) = \int_{\mathcal{M}} P(M|D) \log \frac{P(M|D)}{P(M)} dM. \quad (4)$$

Itti and Baldi (2005) showed that Bayesian surprise predicts attention, as gauged through eye tracking of subjects observing video scenes (Itti and Baldi 2005). Members of the IBM Research group on Computational Creativity developed a recipe generation system, rebranded as *Chef Watson*, which uses Bayesian Surprise to generate recipes that maximize the experience of surprise, controlling for the psychophysics of the recipe’s predicted pleasantness, familiarity and intensity (Sun et al. 2012, Varshney 2013). We suggest that a rich set of backgrounds, experiences and dispositions can be incorporated into the prior to better predict novelty assessments. For example, Murdock et al. (2015) implemented an approximate version of Bayesian Surprise to characterize Charles Darwin’s reading strategy. A topic model trained on his reading up to a given point provides a prior, and the KL divergence between that topic model and one characterizing the next book Darwin read provides a measure of the surprise he might have experienced upon reading it. At different points in his career, Darwin engaged in exploration (seeking out novel texts) and exploitation (digging deeper into familiar topics). More recently, Teplitskiy et al. (2018) draw on peer review data to build an empirical prior for each reviewer. They used reviewers’ training histories, previous coauthors, and published articles to infer the stock of ideas to which reviewers had been exposed. This stock shapes how reviewers cognitively simulate the probability of the manuscripts they evaluate, and hence how they judge novelty.

Refined measurement of surprise and perceived novelty will allow us to assess and predict how people from different eras and locations, with different exposures and commitments, experience a new invention. Consider Figure 6, in which three populations (X , Y , and Z) with distinctive models of the world, experience three new cultural products (A , B , and C). Each population has different expectations, conditioned on their prior knowledge and experience with modes of invention. Product A is unsurprising to population X , but completely jarring to population Z . Product C induces the opposite reaction. To assess this variation empirically, we must be able to enumerate both the prior knowledge familiar to an audience and the customary processes of creation that together condition different subjective perceptions of novelty.

This measurement program would allow us to address important new research questions. How do dispositions, past experiences, and education shape personal simulations of improbability and evaluations of novelty (Bourdieu 1975)? How widely dispersed are novelty assessments because of such differences? Can we characterize and anticipate polarized assessments of a new idea’s novelty, based on these subjective differences? Following work on the experience of suspense and surprise in entertainment, we could also characterize different innovation environments as more or less surprising, in terms of their consistent violation of expectations, and more or less suspenseful, in terms of how new inventions are *expected to disrupt* (Ely et al. 2015).

6. Conclusion

Novelty is central to scientific paradigm shifts, technological revolutions, artistic renaissances, and economic disruptions. Science, technology and cultural policy all rely on measuring novelty, and in some cases intervening to increase or censor it. Yet novelty remains understudied, under-measured, and under-theorized. In this paper, we conceptualize novelty as resting on mental simulations of the search process through which a new invention was created. We argue that these simulations assess novelty as a decreasing function of the probability that an invention *would have emerged* given prior knowledge and the typical inventive process. We then argue that novelty *measures* are a formalization and operationalization of distinctive ideas about the structure of prior knowledge K and the nature of the inventive search process S .

Guided by these insights, we array existing novelty measures along two dimensions: (1) the speed of simulated search S and (2) the complexity of prior knowledge K over which search is simulated. We then summarize these measures with a random walk formalism and show how this expanded view of the discovery process highlights the importance of assessing differences in search speed and complexity in order to accurately measure novelty. To illustrate our argument empirically, we compute several different novelty measures and our random walk approximations on a substantial sample of U.S. patents. We demonstrate that novelty measures rest on different conceptions of the inventive process appropriate to distinct domains of invention. We find that sectors closer in inventive practice have more similar relationships between novelty and impact. We also show that our simple random walk measure justifies our characterization of inventive domains as driven by fast and slow recombinations of components. This empirical demonstration provides striking validation of our conceptual framework, and has clear implications for scholarly practice. Above all, we argue that there is no one novelty measure to rule them all, and that researchers should devote more time and attention to designing or selecting and then justifying the novelty measures they use in their research. Innovation scholars should raise their game in striving for ecologically valid novelty measures.

Building on this framework, we then highlight new directions for novelty measurement that enrich current representations of the search space, including more types of inventive components and black-boxed features that would give analysts a better handle on the multi-scale character of invention. Finally, we show how a Bayesian approach to measuring surprise on the basis of cognitive simulation could support *objectively subjective* measures of perceived novelty in order to account for individual priors based on assumptions and experience. We believe these new directions will allow analysts to make fine-grained predictions of individual variation in novelty assessments. This would be a major step forward in the study of novelty.

We did not focus here on *ex post* measures of innovation or impact, but rather on *ex ante* measurement of novelty. We argue that such measures will increase our ability to analyze and predict the conditions under which a discovery, invention and design becomes important. Here we show that *ex ante* assessments of high novelty tend to be associated with a higher mean and variance in impact (Fleming and Sorenson 2004, Foster et al. 2015).

Much remains to be done in improving assessments of novelty, from improving the data against which novelty is estimated, to enriching the explicit models by which novelty assessment occurs. In an age of increasingly intelligent machines, generative models of novelty can stimulate the creation of novel discoveries, inventions, and designs (Schmidt and Lipson 2009, Sun et al. 2012)—ones that humans will find genuinely surprising and delightful as we refine our ability to measure novelty and predict individual variation in novelty assessment. By providing a unified framework for investigating novelty, and shifting the focus to simulation of the inventive process, we hope to increase understanding of novelty, better predict novelty assessments, and provide templates for the productive analysis, measurement, simulation, and generation of novelty.

Appendix A. Cross-citation and impact profile

Our novelty framework suggests that because different inventive areas have distinctive invention processes, different measures will best characterize novelty in those domains. As novelty is associated with impact, it follows that the relationship between novelty measures and different forms of impact should also vary across inventive areas. To test this, we build three regression models that predict citation counts, hit

Table A.1: Regression results of the cross-citation strength between technology classes on the similarity between impact profiles of novelty measures. Each row presents the coefficient and its p -value of the impact profile similarity in one of the three regression models.

	Coefficient	p -value
Mean Citation	0.6657	0
Hit Citation	0.3430	0
Var Citation	2.6939	0

Table B.2: Test results on the difference between classes for each novelty measure.

	Configurational		Structural		Atypicality		Obscurity		Pecularity	
	statistic	p	statistic	p	statistic	p	statistic	p	statistic	p
ANOVA	2460	0	41	0	591	0	725	0	1019	0
Kruskal	9627	0	419	0	3599	0	4224	0	6766	0
Levene's	45	<.001	87	0	118	<.001	110	<.001	253	<.001
Brown-Forsythe	42	<.001	39	0	116	<.001	109	<.001	213	<.001

patent or not, and the variance of citation, using the five novelty measures implemented above and their squares as predictors. From each model, the estimated coefficients for the novelty measures reveal the differential impact of each implicit novelty-generating process on the reception of patents in that technology class. As a result, we call a particular configuration of coefficients the impact profile of the novelty measures for that class.

We fit our models to each class of patents separately, and test the idea that classes whose inventions frequently cite each other, and presumably have similar inventive practices and patterns, also have similar impact profiles.

Figure A.1 presents the result on the relationship between the cross-citation strength between classes and the impact profile similarity between the classes. Each dot corresponds to a pair of classes, and its x -value is the impact profile similarity between the two classes (i.e., cosine similarity between impact profiles) and its y -value is the cross-citation strength between the classes (i.e., sum of $\log(\text{cross-citation count})$). The impact profiles are derived from the regression models that predicts citation counts (left), hit patent or not (middle), and variance of citations (right). Linear regression lines are fit to the data points, with their 95% confidence intervals. Results from the linear regressions are summarized in Table A.1, where each row presents the coefficient and its p -value of the impact profile similarity in one of the three regression models.

Appendix B. Statistical Tests on Differences in Novelty Distributions between Classes

We perform statistical tests that identify differences in the 5 novelty measures between all the classes. Specifically, we run ANOVA on each novelty measure between the classes to test the difference in the first moments or central tendencies of the measures. As a robustness check, we also perform the Kruskal test (which is non-parametric) on the measures to test the difference in their medians. Similarly, we perform both the Levene's test and the Brown-Forsythe test for the second moments (or variance) of the measures between the classes. The test results are shown in Table B.2. All the test results are significant at $p < .001$, rejecting the null hypothesis that the measures have equal mean (variance) across the classes.

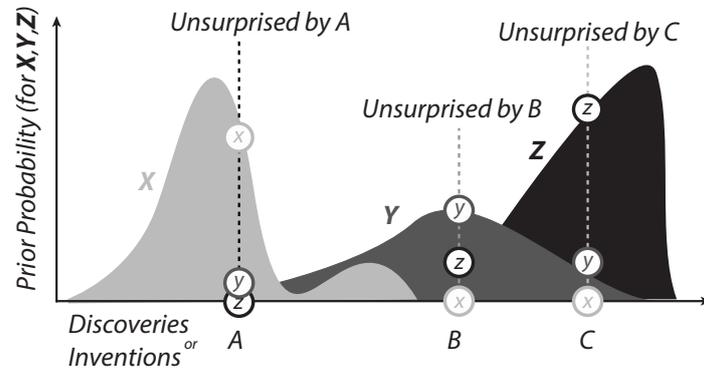


Figure 6: Different levels of experienced surprise based on prior experience and associated expectations.

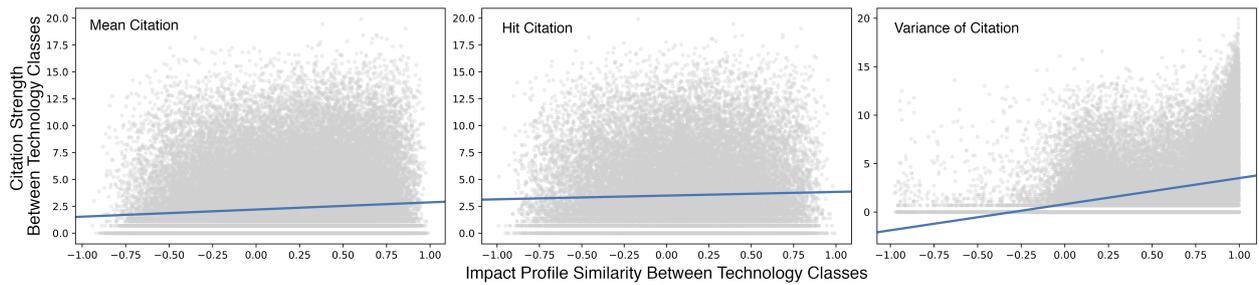


Figure A.1: Regressions of the cross-citation strength between technology classes on the similarity between impact profiles of different novelty measures. Each dot corresponds to a pair of classes, and its x -value is the impact profile similarity between the two classes and its y -value is the cross-citation strength between the classes. The impact profiles are derived from the regression models that predicts logged citation counts (left), hit patent or not (middle), and variance of citations (right). Linear regression lines are fit to the data points, with their 95% confidence intervals.

Appendix C. Descriptive Statistics of the Novelty Measures

Table C.3: Descriptive statistics of the five novelty measures in the eight representative technology classes.

		Configurational	Structural	Atypical	Obscure	Peculiar
504	mean	2.287445	0.236505	-4.368171	-7.075189	0.277460
	std	0.425081	0.490795	1.869919	0.860524	0.120213
	min	1.791759	0.000000	-9.348971	-8.852347	0.116951
	25%	1.945910	0.000000	-5.402069	-7.706878	0.185215
	50%	2.197225	0.000000	-4.013153	-7.073478	0.239460
	75%	2.484907	0.000000	-2.933770	-6.486233	0.353280
	max	4.110874	2.218502	-1.109797	-4.043051	0.770370
	502	mean	2.496725	0.406441	-4.482975	-7.213411
std		0.317273	0.516898	1.278840	0.632452	0.211363
min		1.791759	0.000000	-9.706719	-8.690381	0.144589
25%		2.302585	0.000000	-5.323079	-7.684646	0.370185
50%		2.484907	0.000000	-4.340860	-7.271545	0.459390
75%		2.708050	0.847298	-3.577179	-6.791043	0.597883
max		3.663562	2.069391	-1.097966	-4.892878	1.600453
706		mean	1.963019	0.086629	-6.722817	-5.903123
	std	0.298077	0.252884	1.576375	0.665513	0.177668
	min	1.386294	0.000000	-9.788436	-8.136437	0.299852
	25%	1.791759	0.000000	-8.055577	-6.271797	0.540967
	50%	1.945910	0.000000	-6.939237	-5.913853	0.648114
	75%	2.197225	0.000000	-5.507641	-5.524428	0.757410
	max	3.135494	1.358123	-2.590282	-3.488641	1.588816
	370	mean	1.911232	0.002888	-6.019480	-6.180677
std		0.293463	0.048112	1.247581	0.550496	0.133277
min		1.386294	0.000000	-12.108391	-7.962026	0.245336
25%		1.609438	0.000000	-6.824253	-6.557462	0.505406
50%		1.791759	0.000000	-5.955432	-6.203795	0.587343
75%		2.079442	0.000000	-5.118552	-5.854902	0.674051
max		3.218876	1.308333	-1.128131	-3.828641	1.460020
375		mean	2.217439	0.002143	-5.287776	-6.625595
	std	0.338296	0.038192	1.380575	0.631116	0.175806
	min	1.609438	0.000000	-11.365769	-8.320818	0.173567
	25%	1.945910	0.000000	-6.148246	-7.039203	0.438213
	50%	2.197225	0.000000	-5.085852	-6.627002	0.539636
	75%	2.397895	0.000000	-4.335939	-6.212865	0.666806
	max	3.610918	0.875469	-1.438564	-3.703879	1.472425
	438	mean	2.603157	0.120900	-4.583600	-6.881260
std		0.223442	0.325083	1.204427	0.484095	0.133265
min		2.079442	0.000000	-10.585391	-8.271574	0.152418
25%		2.484907	0.000000	-5.275197	-7.219668	0.315500
50%		2.564949	0.000000	-4.417853	-6.938769	0.371659
75%		2.708050	0.000000	-3.730968	-6.595929	0.440013
max		4.234107	2.124256	-0.949692	-3.748774	1.572956
384		mean	1.840878	0.205614	-7.749334	-5.407451
	std	0.264927	0.330124	1.744041	0.793578	0.372512
	min	1.386294	0.000000	-11.899847	-7.997230	0.263158

	25%	1.609438	0.000000	-8.986353	-5.897154	0.462820
	50%	1.791759	0.000000	-7.794476	-5.411298	0.552571
	75%	1.945910	0.510826	-6.591013	-4.900954	0.721479
	max	2.833213	1.627456	-2.216344	-1.935601	2.740458
	mean	2.317268	0.314295	-6.622708	-5.550741	0.342977
	std	0.267874	0.419101	1.725597	0.711377	0.071758
	min	1.791759	0.000000	-10.668823	-7.546172	0.184106
493	25%	2.079442	0.000000	-7.680559	-6.028714	0.288596
	50%	2.302585	0.000000	-6.799328	-5.470267	0.337079
	75%	2.484907	0.612588	-5.745747	-5.071503	0.386652
	max	3.218876	1.466337	-0.775946	-3.054624	0.580498

Appendix D. Correlation between Novel Measures

Table D.4: Pearson correlation between the five novelty measures in each of the eight representative technology classes.

	Configurational	Structural	Atypical	Obscure	Peculiar
504					
Configurational	1.000000	0.428182	0.558573	-0.245412	-0.318711
Structural	0.428182	1.000000	0.180166	0.173319	0.118111
Atypical	0.558573	0.180166	1.000000	-0.639300	-0.466158
Obscure	-0.245412	0.173319	-0.639300	1.000000	0.356249
Peculiar	-0.318711	0.118111	-0.466158	0.356249	1.000000
502					
Configurational	1.000000	0.394207	0.496632	-0.106152	-0.146387
Structural	0.394207	1.000000	0.339393	0.022938	-0.381984
Atypical	0.496632	0.339393	1.000000	-0.379368	-0.133533
Obscure	-0.106152	0.022938	-0.379368	1.000000	-0.518366
Peculiar	-0.146387	-0.381984	-0.133533	-0.518366	1.000000
706					
Configurational	1.000000	0.282647	0.491938	-0.099658	-0.386569
Structural	0.282647	1.000000	0.244430	-0.056188	-0.122171
Atypical	0.491938	0.244430	1.000000	-0.539642	-0.145569
Obscure	-0.099658	-0.056188	-0.539642	1.000000	-0.301380
Peculiar	-0.386569	-0.122171	-0.145569	-0.301380	1.000000
370					
Configurational	1.000000	0.076071	0.499094	-0.036527	-0.261471
Structural	0.076071	1.000000	0.007449	0.039373	-0.046979
Atypical	0.499094	0.007449	1.000000	-0.492275	-0.179773
Obscure	-0.036527	0.039373	-0.492275	1.000000	-0.366612
Peculiar	-0.261471	-0.046979	-0.179773	-0.366612	1.000000
375					
Configurational	1.000000	0.006581	0.534028	-0.175022	-0.405630
Structural	0.006581	1.000000	-0.014529	0.029358	-0.024789
Atypical	0.534028	-0.014529	1.000000	-0.488141	-0.318978
Obscure	-0.175022	0.029358	-0.488141	1.000000	-0.242031
Peculiar	-0.405630	-0.024789	-0.318978	-0.242031	1.000000
438					
Configurational	1.000000	0.211768	0.530743	-0.131935	-0.111380
Structural	0.211768	1.000000	0.171697	0.008742	-0.027515
Atypical	0.530743	0.171697	1.000000	-0.437946	-0.170901
Obscure	-0.131935	0.008742	-0.437946	1.000000	-0.448252
Peculiar	-0.111380	-0.027515	-0.170901	-0.448252	1.000000
384					
Configurational	1.000000	0.501964	0.549510	-0.189059	-0.393557
Structural	0.501964	1.000000	0.485803	-0.358065	-0.185276
Atypical	0.549510	0.485803	1.000000	-0.679585	-0.394297
Obscure	-0.189059	-0.358065	-0.679585	1.000000	-0.039395
Peculiar	-0.393557	-0.185276	-0.394297	-0.039395	1.000000
493					
Configurational	1.000000	0.430825	0.663964	-0.365170	-0.198448
Structural	0.430825	1.000000	0.241609	-0.102010	-0.016584

Atypical	0.663964	0.241609	1.000000	-0.734429	-0.157017
Obscure	-0.365170	-0.102010	-0.734429	1.000000	-0.172684
Peculiar	-0.198448	-0.016584	-0.157017	-0.172684	1.000000

Appendix E. Results for Patents Granted in 1990

We focus on patents granted in 2000 in Section 4; here we confirm our findings using 90,421 utility patents granted in 1990. First, Figure E.2 confirms that patents from different technology areas vary substantially in their distribution of novelty scores across measures (cf. Figure 2). All the statistical tests (ANOVA, Kruskal, Levene, and Brown-Forsythe) are significant at .001 level, supporting that the novelty distributions are different across the classes. We further show in Figure 3 that fields where patents most intensively cite one another—and presumably share inventive processes—have greater similarity in impact profiles. This is consistent with our claim that in different technology areas, distinct measures are important to the prediction of impact (mean, hit, and variance). Finally, in Figure 5, we show how mutual information between the random walk novelty and every other novelty measure varies as the walk speed increases. For configurational and structural novelty, mutual information decreases as random walk speed increases, with configurational novelty dropping the fastest ($\beta = -0.0069$, $p < .001$) and structural novelty dropping gradually ($\beta = -0.0023$, $p < .001$). For combinatorial novelties, mutual information increases as random walk speed increases ($\beta_{peculiarity} = 0.007$, $p < .001$; $\beta_{atypicality} = 0.0309$, $p < .001$; $\beta_{obscurity} = 0.0186$, $p < .001$). This pattern confirms our conceptual framework that configurational novelty embodies a slow, local search process while combinatorial novelty assumes a fast, non-local one.

References

- Arthur, B. 2009. The nature of technology. *New York: Free Press* .
- Austen, Jane. 1813. *Pride and prejudice*. printed for T. Egerton, Military Library, Whitehall.
- Austen, Jane, Seth Grahame-Smith. 2009. *Pride and prejudice and zombies*. Quirk Books.
- Azoulay, Pierre, Christopher C Liu, Toby E Stuart. 2017. Social influence given (partially) deliberate matching: Career imprints in the creation of academic entrepreneurs. *American Journal of Sociology* **122**(4) 1223–1271.
- Barron, Alexander TJ, Jenny Huang, Rebecca L Spang, Simon DeDeo. 2018. Individuals, institutions, and innovation in the debates of the french revolution. *Proceedings of the National Academy of Sciences* **115**(18) 4607–4612.
- Bengio, Yoshua, Aaron Courville, Pascal Vincent. 2013. Representation learning: A review and new perspectives. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **35**(8) 1798–1828.
- Benner, Mary J., Michael L. Tushman. 2003. Exploitation, exploration, and process management: The productivity dilemma revisited. *Academy of Management Review* **28**(2) 238–256. doi:10.5465/AMR.2003.9416096.
- Blei, David M., Andrew Y. Ng, Michael I. Jordan. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.* **3** 993–1022.
- Blondel, Vincent D, Jean-Loup Guillaume, Renaud Lambiotte, Etienne Lefebvre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* **2008**(10) P10008. doi: 10.1088/1742-5468/2008/10/p10008.
- Blundell, Richard, Rachel Griffith, John Van Reenen. 1999. Market share, market value and innovation in a panel of british manufacturing firms. *The review of economic studies* **66**(3) 529–554.
- Bonacich, Phillip. 1987. Power and centrality: A family of measures. *American Journal of Sociology* **92**(5) 1170–1182.
- Bonacich, Phillip, Paulette Lloyd. 2001. Eigenvector-like measures of centrality for asymmetric relations. *Social Networks* **23**(3) 191–201. doi:10.1016/S0378-8733(01)00038-7.
- Bourdieu, P. 1975. The specificity of the scientific field and the social conditions for the progress of reason. *Soc. Sci. Inf.* **14** 19–47.
- Burt, R. S. 2004. Structural holes and good ideas. *American Journal of Sociology* **110**(2) 349–399.
- Carnabuci, Gianluca, Jeroen Bruggeman. 2009. Knowledge specialization, knowledge brokerage and the uneven growth of technology domains. *Social Forces* **88**(2) 607–641. doi:10.1353/sof.0.0257.
- Chen, C., Y. Chen, M. Horowitz, H. Hou, Z. Liu, D. Pellegrino. 2009. Towards an explanatory and computation theory of scientific discovery. *Journal of Informetrics* **3** 191–209.

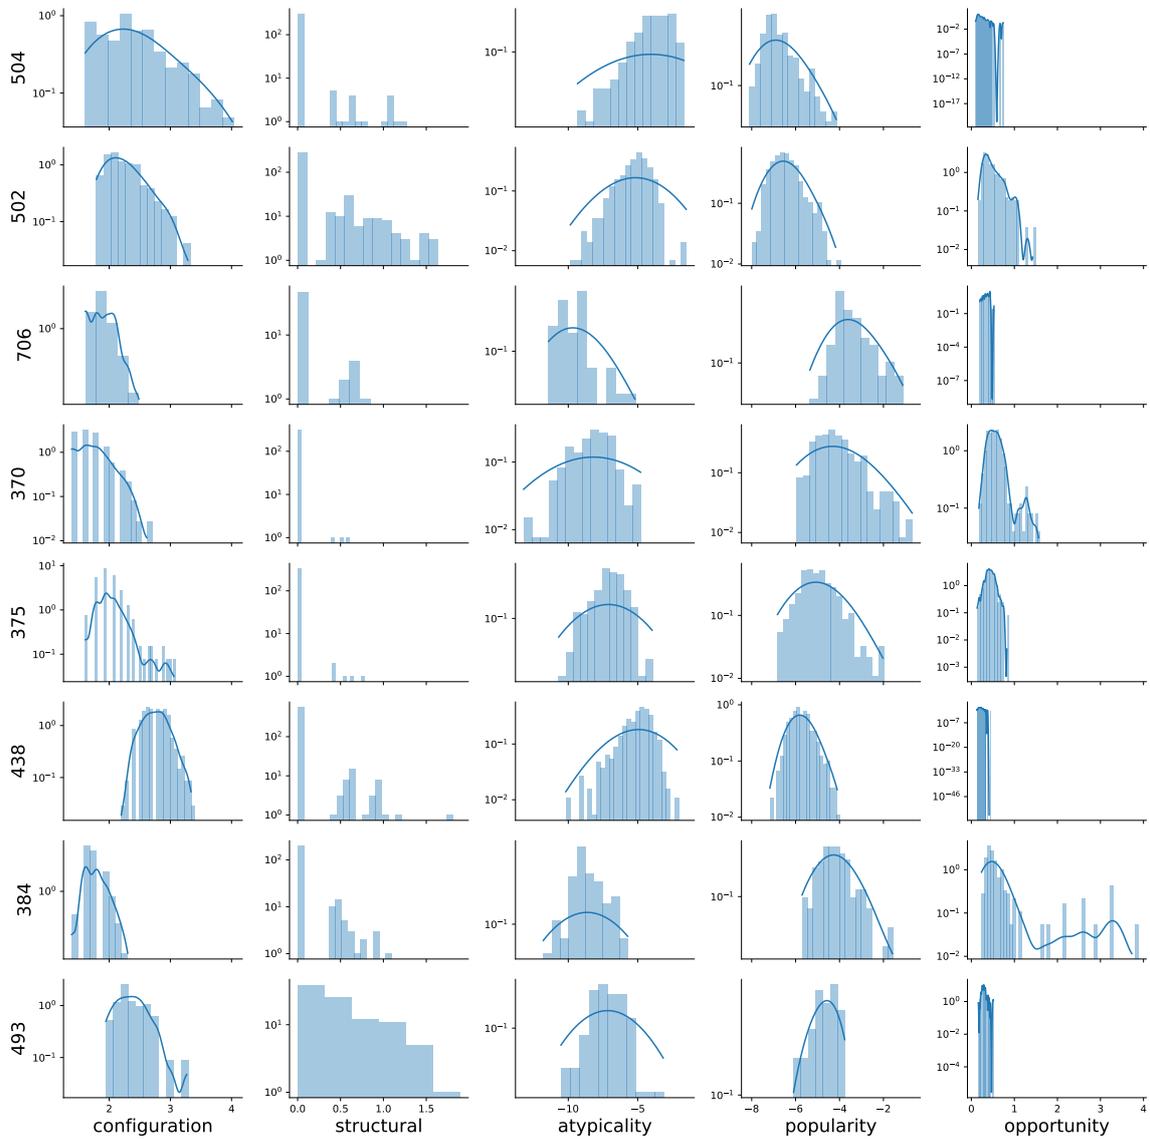


Figure E.2: Distinct technology classes (y-axis labels) from different USPTO “technology centers” reveal marked differences in the distribution of novelty measures. The novelty scores are computed for patents granted in 1990. Solid lines are kernel density estimates with Gaussian kernels.

- Cheng, X-Q., F-X. Ren, H-W. Shen, Z-K. Zhang, T. Zhao. 2010. Bridgeness: a local index on edge significance in maintaining global connectivity. *Journal of Statistical Mechanics: Theory and Experiment* **2010**. 10011.
- Christensen, Clayton M. 2011. *The Innovator's Dilemma: The Revolutionary Book That Will Change the Way You Do Business*. Reprint edition ed. HarperBusiness.
- Clark, Andy. 2015. *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.
- Cohen, Wesley M. 2010. Fifty years of empirical studies of innovative activity and performance. *Handbook of the Economics of Innovation*, vol. 1. Elsevier, 129–213.
- Cohen, Wesley M, You-Na Lee, John P Walsh. 2019. How innovative are innovations? a multidimensional, survey-based approach. *Measuring and Accounting for Innovation in the 21st Century*. University of Chicago Press.
- Cokol, M., I. Iossifov, C. Weinreb, A. Rzhetsky. 2005. Emergent behavior of growing knowledge about molecular interactions. *Nat Biotechnol* **23**(10) 1243–7.
- Coombs, Rod, P Narandren, A Richards. 1996. A literature-based innovation output indicator. *Research policy* **25**(3) 403–413.
- Cowell, Frank A. 2000. Measurement of inequality. *Handbook of income distribution* **1** 87–166.
- Decelle, Aurelien, Florent Krzakala, Cristopher Moore, Lenka Zdeborová. 2011. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E* **84**(6) 066106.
- Deerwester, Scott C., Susan T Dumais, Thomas K. Landauer, George W. Furnas, Richard A. Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American Society for Information Science* **41**(6) 391–407.
- Delort, Jean-Yves, Enrique Alfonseca. 2012. DualSum: a topic-model based approach for update summarization. *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 214–223.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* .
- Dubois, Pierre, Olivier De Mouzon, Fiona Scott-Morton, Paul Seabright. 2015. Market size and pharmaceutical innovation. *The RAND Journal of Economics* **46**(4) 844–871.
- Ely, Jeffrey, Alexander Frankel, Emir Kamenica. 2015. Suspense and surprise. *Journal of Political Economy* **123**(1) 215–260.
- Erwin, Douglas H. 2020. A conceptual framework of evolutionary novelty and innovation. *Biological Reviews* .
- Evans, James A. 2010. Industry induces academic science to know less about more. *American Journal of Sociology* **116**(2) 389–452. doi:10.1086/653834.
- Fleming, L., O. Sorenson. 2001. Technology as a complex adaptive system: evidence from patent data. *Research Policy* **30**(7) 1019–1039.
- Fleming, L., O. Sorenson. 2004. Science as a map in technological search. *Strategic Management Journal* **25**(8) 909–928. doi:10.1002/Smj.384.
- Fleming, Lee. 2001. Recombinant uncertainty in technological search. *Management science* **47**(1) 117–132.
- Fontana, Magda, Martina Iori, Fabio Montobbio, Roberta Sinatra. 2020. New and atypical combinations: An assessment of novelty and interdisciplinarity. *Research Policy* **49**(7) 104063.
- Fortunato, Santo, Darko Hric. 2016. Community detection in networks: A user guide. *Physics Reports* **659** 1–44.
- Fosfuri, Andrea, Marco S Giarratana. 2009. Masters of war: Rivals' product innovation and new advertising in mature product markets. *Management Science* **55**(2) 181–191.
- Foster, David V., Mary M. Rorick, Tanja Gesell, Laura M. Feeney, Jacob G. Foster. 2013a. Dynamic landscapes: A model of context and contingency in evolution. *Journal of Theoretical Biology* **334** 162–172. doi: 10.1016/j.jtbi.2013.05.030.
- Foster, David V, Mary M Rorick, Tanja Gesell, Laura M Feeney, Jacob G Foster. 2013b. Dynamic landscapes: A model of context and contingency in evolution. *Journal of theoretical biology* **334** 162–172.
- Foster, Jacob G., Andrey Rzhetsky, James A. Evans. 2015. Tradition and innovation in scientists research strategies. *American Sociological Review* **80**(5) 875–908. doi:10.1177/0003122415601618.
- Foster, J.G., D.V. Foster, P. Grassberger, M. Paczuski. 2007. Link and subgraph likelihoods in random undirected networks with fixed and partially fixed degree sequences. *Phys. Rev. E* **76** 046112.

- Funk, Russell J, Jason Owen-Smith. 2016. A dynamic network measure of technological change. *Management Science* **63**(3) 791–817.
- Galloway, Alexander R, Eugene Thacker. 2007. *The exploit: A theory of networks*, vol. 21. University of Minnesota Press.
- Garg, Nikhil, Londa Schiebinger, Dan Jurafsky, James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* **115**(16) E3635–E3644.
- Gavrilets, Sergey. 1999. Dynamics of clade diversification on the morphological hypercube. *Proceedings of the Royal Society of London. Series B: Biological Sciences* **266**(1421) 817–824.
- Gell, Alfred. 1998. *Art and agency: an anthropological theory*. Clarendon Press.
- Geroski, Paul A, Richard Pomroy. 1990. Innovation and the evolution of market structure. *The journal of industrial economics* 299–314.
- Gerow, Aaron, Yuening Hu, Jordan Boyd-Graber, David M Blei, James A Evans. 2018. Measuring discursive influence across scholarship. *Proceedings of the national academy of sciences* **115**(13) 3308–3313.
- Gerrish, Sean, David M. Blei. 2010. A language-based approach to measuring scholarly impact. *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*. 375–382.
- Goodfellow, Ian, Yoshua Bengio, Aaron Courville, Yoshua Bengio. 2016. *Deep learning*, vol. 1. MIT press Cambridge.
- Gross, Daniel P. 2014. Creativity under fire: The effects of competition on innovation and the creative process. Available at SSRN 2520123 .
- Grossman, Robert L., Michael Sabala, Yunhong Gu, Anushka Anand, Matt Handley, Rajmonda Sulo, Lee Wilkinson. 2008. *Discovering Emergent Behavior From Network Packet Data: Lessons from the Angle Project*.
- Guinan, Eva, Kevin J Boudreau, Karim R Lakhani. 2013. Experiments in open innovation at harvard medical school. *MIT Sloan Management Review* **54**(3) 45.
- Harhoff, Dietmar, Francis Narin, Frederic M Scherer, Katrin Vopel. 1999. Citation frequency and the value of patented inventions. *Review of Economics and statistics* **81**(3) 511–515.
- Henderson, Rebecca M, Kim B Clark. 1990. Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Administrative science quarterly* 9–30.
- Hutchens, Robert. 2001. Numerical measures of segregation: desirable properties and their implications. *Mathematical social sciences* **42**(1) 13–29.
- Itti, Laurent, Pierre F. Baldi. 2005. Bayesian surprise attracts human attention. *Advances in neural information processing systems*. 547–554.
- Jaffe, Adam B, Manuel Trajtenberg, Rebecca Henderson. 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *the Quarterly journal of Economics* **108**(3) 577–598.
- Johnson, Steven. 2011. *Where Good Ideas Come From*. Reprint edition ed. Riverhead Trade.
- Joulin, Armand, Edouard Grave, Piotr Bojanowski, Tomas Mikolov. 2016. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759* .
- Kauffman, Stuart A. 2000. *Investigations*. Oxford University Press.
- Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, Matt Taddy. under review. Measuring technological innovation over the long run. *preprint* .
- Kemp, René, Peter Pearson. 2008. Final report of the project measuring eco-innovation. *Maastricht (The Netherlands)* .
- Klingenstein, Sara, Tim Hitchcock, Simon DeDeo. 2014. The civilizing process in londons old bailey. *Proceedings of the National Academy of Sciences* **111**(26) 9419–9424. doi:10.1073/pnas.1405984111.
- Kozlowski, Austin C, Matt Taddy, James A Evans. 2019. The geometry of culture: Analyzing the meanings of class through word embeddings. *American Sociological Review* **84**(5) 905–949.
- Kuhn, Thomas. 1962. *The Structure of Scientific Revolutions*. University of Chicago Press.
- Lamont, Michle. 2009. *How professors think : inside the curious world of academic judgment*. Harvard University Press.
- Langville, Amy N., C. D. Meyer. 2006. *Google’s PageRank and beyond : the science of search engine rankings*. Princeton University Press.

- Larkey, Leah S., James Allan, Margaret E. Connell, Alvaro Bolivar, Courtney Wade. 2002. UMass at TREC 2002: Cross language and novelty tracks.
- Latour, Bruno. 1987. *Science in action : how to follow scientists and engineers through society*. Harvard University Press.
- Leahey, Erin, Christine Beckman, Taryn Stanko. 2015. Prominent but less productive: The impact of interdisciplinarity on scientists' research. *arXiv:1510.06802 [cs]* .
- Leahey, Erin, James Moody. 2014. Sociological innovation through subfield integration. *Social Currents* **1**(3) 228–256. doi:10.1177/2329496514540131.
- Lee, Jinhyuk, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* **36**(4) 1234–1240.
- Lee, You-Na. 2015. Evaluating and extending innovation indicators for innovation policy. *Research Evaluation* **24**(4) 471–488.
- Lee, You-Na, John P Walsh, Jian Wang. 2015. Creativity in scientific teams: Unpacking novelty and impact. *Research Policy* **44**(3) 684–697.
- Levy, Omer, Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. *Advances in neural information processing systems* **27** 2177–2185.
- Lin, Yiling, James Allen Evans, Lingfei Wu. 2021. Novelty, disruption, and the evolution of scientific impact. *arXiv preprint arXiv:2103.03398* .
- Liu, Meijun, Yi Bu, Chongyan Chen, Jian Xu, Daifeng Li, Yan Leng, Richard Barry Freeman, Eric Meyer, Wonjin Yoon, Mujeen Sung, et al. 2020. Can pandemics transform scientific novelty? evidence from covid-19. *arXiv preprint arXiv:2009.12500* .
- Margolis, Howard. 1993. *Paradigms and Barriers: How Habits of Mind Govern Scientific Beliefs*. 1st ed. University Of Chicago Press.
- Mawer, Deborah. 2000. *The Cambridge Companion to Ravel*. Cambridge University Press.
- Mikolov, Tomas, Kai Chen, Greg Corrado, Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* .
- Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S Corrado, Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*. 3111–3119.
- Montroll, Elliott W., Harvey Scher. 1973. Random walks on lattices. IV. continuous-time walks and influence of absorbing boundaries. *Journal of Statistical Physics* **9**(2) 101–135. doi:10.1007/BF01016843.
- Moser, Petra. 2005. How do patent laws influence innovation? evidence from nineteenth-century world's fairs. *American economic review* **95**(4) 1214–1236.
- Murdock, Jaimie, Colin Allen, Simon DeDeo. 2015. Exploration and exploitation of victorian science in darwin's reading notebooks. *arXiv preprint arXiv:1509.07175* .
- Nanda, Ramana, Ken Younge, Lee Fleming. 2013. Innovation and entrepreneurship in renewable energy. *The Changing Frontier: Rethinking Science and Innovation Policy*. University of Chicago Press.
- Newman, M. 2003. The structure and function of complex networks. *SIAM Review* **45**(2) 167–256. doi:10.1137/S003614450342480.
- Nichols, Leah G. 2014. A topic model approach to measuring interdisciplinarity at the national science foundation. *Scientometrics* **100**(3) 741–754. doi:10.1007/s11192-014-1319-2.
- North, Michael. 2013. *Novelty: A history of the new*. University of Chicago Press.
- OED, 7th edition. 2007. “novel”. Angus Stevenson, ed., *The shorter oxford english dictionary: on historical principles*, 7th ed. Oxford: Clarendon Press.
- Padgett, John, Chris Ansell. 1992. Robust action and the rise of the medici. *American Journal of Sociology* **98** 1259–1330.
- Padgett, John F., Walter W. Powell. 2012. *The Emergence of Organizations and Markets*. Princeton University Press.
- Pennington, Jeffrey, Richard Socher, Christopher Manning. 2014. Glove: Global vectors for word representation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 1532–1543.

- Poincaré, Henri. 1952. Mathematical creation. Brewster Ghiselin, ed., *The Creative Process: A Symposium*. University of California Press.
- Pontikes, Elizabeth. 2014. Perspectives on recombination.
- Powell, W. W., K. W. Koput, L. SmithDoerr. 1996. Interorganizational collaboration and the locus of innovation: Networks of learning in biotechnology. *Administrative Science Quarterly* **41**(1) 116–145.
- Rohrmeier, Martin A, Stefan Koelsch. 2012. Predictive information processing in music cognition. a critical review. *International Journal of Psychophysiology* **83**(2) 164–175.
- Romer, Paul M. 1986. Increasing returns and long-run growth. *Journal of political economy* **94**(5) 1002–1037.
- Romer, Paul M. 1994. The origins of endogenous growth. *Journal of Economic perspectives* **8**(1) 3–22.
- Ross, Alex. 2007. *The rest is noise: Listening to the twentieth century*. Macmillan.
- Rosvall, Martin, Carl T Bergstrom. 2010. Mapping change in large networks. *PLoS one* **5**(1) e8694.
- Rzhetsky, Andrey, Jacob G. Foster, Ian T. Foster, James A. Evans. 2015. Choosing experiments to accelerate collective discovery. *Proc Natl Acad Sci U S A* **112**(47) 14569–14574. doi:10.1073/pnas.1509757112.
- Schilling, Melissa A, Elad Green. 2011. Recombinant search and breakthrough idea generation: An analysis of high impact papers in the social sciences. *Research Policy* **40**(10) 1321–1331.
- Schmidt, M., H. Lipson. 2009. Distilling free-form natural laws from experimental data. *Science* **324**(5923) 81–5. doi:10.1126/science.1165893.
- Seneviratne, Dilesha, Shlomo Geva, Guido Zuccon, Gabriela Ferraro, Timothy Chappell, Magali Meireles. 2015. A signature approach to patent classification. *AIRS*. Springer, 413–419.
- Shi, Feng, Jacob G. Foster, James A. Evans. 2015. Weaving the fabric of science: Dynamic network models of science’s unfolding structure. *Social Networks* **43** 73–85. doi:10.1016/j.socnet.2015.02.006. 00001.
- Simmons, Matthew P., Lada A. Adamic, Eytan Adar. 2011. Memes online: Extracted, subtracted, injected, and recollected. *ICWSM* **11** 17–21.
- Simon, Herbert A. 1969. The sciences of the artificial. *Cambridge, MA* .
- Sun, John Z., Grace I. Wang, Vivek K Goyal, Lav R. Varshney. 2012. A framework for bayesian optimality of psychophysical laws. *Journal of Mathematical Psychology* **56**(6) 495–501. doi:10.1016/j.jmp.2012.08.002.
- Teplitskiy, Misha, Daniel Acuna, Aïda Elamrani-Raoult, Konrad Körding, James Evans. 2018. The sociology of scientific validity: How professional networks shape judgement in peer review. *Research Policy* **47**(9) 1825–1841.
- Tria, F., V. Loreto, V. D. P. Servedio, S. H. Strogatz. 2014. The dynamics of correlated novelties. *Scientific Reports* **4**. doi:10.1038/srep05890. 00003.
- Tribus, Myron. 1961. *Thermostatistics and thermodynamics*. New York : Van Nostrand.
- Tshitoyan, Vahe, John Dagdelen, Leigh Weston, Alexander Dunn, Ziqin Rong, Olga Kononova, Kristin A Persson, Gerbrand Ceder, Anubhav Jain. 2019. Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature* **571**(7763) 95–98.
- Tushman, Michael, David Nadler, et al. 1986. Organizing for innovation. *California management review* **28**(3) 74–92.
- Tversky, Amos, Daniel Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *science* **185**(4157) 1124–1131.
- Uzzi, Brian, Satyam Mukherjee, Michael Stringer, Ben Jones. 2013. Atypical combinations and scientific impact. *Science* **342**(6157) 468–472. doi:10.1126/science.1240474.
- Varshney, L.R. 2013. To surprise and inform. *2013 IEEE International Symposium on Information Theory Proceedings (ISIT)*. 3145–3149. doi:10.1109/ISIT.2013.6620805.
- Verhoeven, Dennis, Jurriën Bakker, Reinhilde Veugelers. 2016. Measuring technological novelty with patent-based indicators. *Research Policy* **45**(3) 707–723.
- Wagner, Andreas, William Rosen. 2014. Spaces of the possible: universal darwinism and the wall between technological and biological innovation. *Journal of The Royal Society Interface* **11**(97) 20131190.
- Wang, Dashun, Chaoming Song, Albert-Lszl Barabasi. 2013. Quantifying long-term scientific impact. *Science* **342**(6154) 127–132. doi:10.1126/science.1237825.
- Wu, Lingfei, Dashun Wang, James A Evans. 2019. Large teams develop and small teams disrupt science and technology. *Nature* **566**(7744) 378–382.

Yan, Bowen, Jianxi Luo. 2017. Measuring technological distance for patent mapping. *Journal of the Association for Information Science and Technology* **68**(2) 423–437.