

Aspects of theory-ladenness in data-intensive science¹

Wolfgang Pietsch²
Munich Center for Technology in Society,
TU München, Munich, Germany

Abstract

Recent claims, mainly from computer scientists, concerning a largely automated and model-free data-intensive science have been countered by critical reactions from a number of philosophers of science. The debate suffers from a lack of detail in two respects, regarding (i) the actual methods used in data-intensive science and (ii) the specific ways in which these methods presuppose theoretical assumptions. I examine two widely-used algorithms, classificatory trees and non-parametric regression, and argue that these are theory-laden in an external sense, regarding the framing of research questions, but not in an internal sense concerning the causal structure of the examined phenomenon. With respect to the novelty of data-intensive science, I draw an analogy to exploratory as opposed to theory-directed experimentation.

1. Introduction

Over the past decade, computer scientists have claimed that a new scientific methodology has become possible through advances in information technology (e.g. Gray 2007). This approach is supposed to be data-driven, strongly inductive, and relatively theory-independent. The epistemology of such data-intensive science has recently emerged as a novel topic in philosophy of science. Generally, the reactions have been rather critical, often referring to the more or less trivial observation that some kind of theory-ladenness always occurs in scientific research. But, as I will argue, this means throwing the baby out with the bathwater, since interesting shifts in the role of theory can indeed be observed when examining specific methods employed in data-intensive science.

In Section 2, I will suggest a definition for data-intensive science reflecting those features that are interesting from an epistemological perspective. I will then, in Section 3, briefly introduce the debate on theory-ladenness in data-intensive science. To assess the various arguments, I will discuss two algorithms that are widely used, namely classificatory trees (Section 4) and non-parametric regression (Section 5). For both of these methods, I will point out the specific ways in which theory has to be presupposed to identify causal connections and thus yield reliable predictions. I will conclude in Section 6 that these algorithms require an *external* theory-ladenness concerning the framing of research questions, but little *internal* theory-ladenness concerning the causal structure of the examined phenomena. I will also point out remarkable analogies to the analysis of theory-ladenness in exploratory experimentation.

¹ accepted for the proceedings volume of PSA2014, 24th Biennial Meeting of the Philosophy of Science Association

² pietsch@cvi-a.tum.de

2. Defining data-intensive science

The problems usually addressed in data-intensive science bear close resemblance to standard problems in statistics. They concern classification or regression of an output variable y with respect to a large number of input parameters x , also called predictor variables or covariates, on the basis of large training sets. The main differences compared with conventional problems in statistics consist in the high-dimensionality of the input variable and the amount of data available about various configurations or states of the system. For example, an internet store wants to know how likely someone buys a certain product depending on surf history, various cookies and a user profile as well as based on data of other users who have either bought or failed to buy the product. A medical researcher examines which combinations of genetic and environmental factors are responsible for a certain disease. A political adviser is interested how likely a specific individual is going to vote for a certain candidate based on a profile combining for example voting history, political opinions, general demographics, or consumer data.

In a *classification* problem, the output variable has a finite number of possible values. In a *regression* problem, the output variable is continuous. In order to establish an adequate and reliable model, extensive training and test data is needed. Each instance in the training and test sets gives a value for the output variable dependent on at least some of the input parameters. The training data is used to build the model, e.g. determine relevant parameters, the test data to validate and verify the model. Using part of the data to determine the accuracy of a model is commonly referred to as cross-validation.³

In this essay, we cannot delve into all the technical details of the various algorithms employed in data-intensive science, such as support vector machines, forests or neural networks. Instead we will look at two simple but widely-used algorithms, namely classificatory trees and non-parametric regression, to examine how much and what kind of theory must be presupposed in order for these algorithms to yield meaningful results.

The term *data-intensive science* is notoriously blurry, as has been emphasized for example by Sabina Leonelli: ‘a general characterisation of data-driven methods is hard to achieve, given the wide range of activities and epistemic goals currently subsumed under this heading.’ (2012, 1) However, in order to say something substantial about the role of theory, we have to be more specific about the kinds of practices we want to include as data-intensive science even if an exact definition does not fully correspond to common usage of the term.

In the computer science literature, various definitions have been proposed for the closely related concepts of a *data deluge* or of *big data*. Most of these refer to the pure amount of information or to the technical challenges that such ‘big data’ poses in terms of the so-called ‘three Vs’—volume, velocity and variety of data (Laney 2001). However, from a philosophy of science perspective, these definitions do not provide much insight. After all, larger amounts of data do not automatically imply interesting methodological developments.

³ An excellent introductory textbook is Russell & Norvig (2009).

Leonelli, partly following Gray (2007, xix), identifies two characteristic features for data-intensive methodology: ‘one is the intuition that induction from existing data is being vindicated as a crucial form of scientific inference, which can guide and inform experimental research; and the other is the central role of machines, and thus of automated reasoning, in extracting meaningful patterns from data.’ (2012, 1) She adds that these features are themselves quite controversial and criticizes that they are difficult to apply in research contexts.

In defining data-intensive science, I largely follow Leonelli, while attempting to be more precise about the type of induction. I will argue that eliminative induction in the tradition of Mill’s methods⁴ plays the crucial role. The first part of my definition thus focusses on the premises that are necessary to carry out eliminative induction: data-intensive science requires (I) *data representing all relevant configurations of the examined phenomenon with respect to a specific research question*. For complex phenomena, this implies high-dimensional data, i.e. data sets involving many parameters, as well as a large number of observations or instances covering a wide range of combinations of these parameters. We will see later that this premise underwrites the characteristic data-driven and inductive nature of data-intensive science.

(II) The second feature concerns the *automation of the entire scientific process*, from data capture to processing to modeling (cp. Gray 2007, xix). This allows sidestepping some of the limitations of the human cognitive apparatus but also leads to a loss in human understanding regarding the results of data-intensive science. Again, being more precise about the type of induction allows to determine under which circumstances automation is really possible.

3. Theory-free science?

Proponents of data-intensive science claim that important changes are happening with respect to the role of theory. An extreme, but highly influential version of such a statement is by the former editor-in-chief of *Wired* Chris Anderson, who notoriously proclaimed ‘the end of theory’ altogether (2008). More nuanced positions can be found for example in the writings of Google research director Peter Norvig (2009): ‘Having more data, and more ways to process it, means that we can develop different kinds of theories and models.’ Simpler models with a lot of data supposedly trump more elaborate models with less data (Halevy et al. 2009, 9).

A number of philosophers have objected to claims of a theory-free science—generally by pointing out various kinds of theory-ladenness. For example, Werner Callebaut writes: ‘We know from Kuhn, Feyerabend, and [...] Popper that observations (facts, data) are theory-laden. Popper [...] rejected the “bucket theory of knowledge” in favor of the “searchlight theory,” according to which observation “is a process in which we play an intensely active part.” Our perceptions are always preceded by interests, questions, or expectations—in short, by something “speculative”.’ (Callebaut 2012, 74) Leonelli concurs in her work on big data biology: ‘Using data for the purposes of discovery can happen in a variety of ways, and involves a complex ensemble of skills and methodological components. Inferential reasoning

⁴ not to be confused with a looser use of the same term in the sense of eliminating hypotheses until only the correct one remains

from data is tightly interrelated with specific theoretical commitments about the nature of the biological phenomena under investigation, as well as with experimental practices through which data are produced, tested and modelled. For instance, extracting biologically meaningful inferences from high-throughput genomic data may involve reliance on theories about gene expression and regulation, models of the biological processes being regulated and familiarity with the instruments and organisms from which data were obtained. In this context, “inductive” clearly does not mean “hypothesis-free”; nor can automated reasoning be seen as a substitute to human judgment based on specific expertise and laboratory experience.’ (2012, 2)

Certainly, the idea of an entirely theory- or model-free science is absurd. So, Callebaut and Leonelli rightly point out various kinds of theoretical assumptions that enter scientific analyses. But this kind of argument turns out too general and in the end fails to do justice to the remarkable shift towards a strongly inductive approach. Thus, the interesting question is in which ways data-intensive science is indeed theory-laden, and, more importantly, in which sense it can be theory-free. To provide an answer, we now take a detailed look at two algorithms that are widely employed, namely classificatory trees and non-parametric regression. We link these methods to eliminative induction and then determine the kind of theoretical knowledge that has to be presupposed.

4. First case study: classificatory trees

Classificatory trees (e.g. Russell & Norvig 2010, Ch. 18.3.3) are used to determine whether a certain instance belongs to a particular group A depending on a number of parameters C_1, \dots, C_N and thus perfectly match the scheme of data-intensive problems as described in Section 2. With help of training data, the tree is set up recursively. First, the parameter C_X is determined that contains the largest amount of information with respect to the classification of the training data, as formally measured in terms of Shannon entropy. If C_X classifies all instances correctly, the procedure is terminated. Otherwise, two subproblems remain, namely classifying when C_X is present and when it is absent. This step is repeated until either all instances are classified correctly or no potential classifiers are left. If the algorithm is successful, the resulting tree structure gives a Boolean expression of necessary and sufficient conditions for A , which can be interpreted as a complex scientific law: e.g. if $(C_3C_2 \vee C_4\neg C_2)C_1 \vee C_6C_5\neg C_1$, then A .

The framing of classificatory trees in particular and of problems in data-intensive science in general in terms of a mapping of boundary conditions to an outcome variable fits well with eliminative induction as exemplified in John Stuart Mill’s *methods of elimination* (1886, Bk. III, Ch. VIII) with a predecessor in Francis Bacon’s *method of exclusion* (1620/1994, Bk. 2). While until the end of the 19th century, Bacon’s approach was widely considered the methodological foundation for modern science, eliminative induction has not been very popular since. So, there exist comparably few modern accounts, including von Wright (1951), Mackie (1965, appendix), Skyrms (2000), Baumgartner & Grasshoff (2004), Pietsch (2014).⁵

⁵ In the following, I will largely rely on the last account.

In eliminative induction, a phenomenon *A* is examined under the systematic variation of potentially relevant boundary conditions *C*₁, ..., *C*_{*N*} with the aim of establishing *causal relevance* or *irrelevance* of these conditions, relative to a certain context or background *B* consisting of further boundary conditions. The best known and arguably most effective method is the so-called *method of difference* that establishes causal relevance of a boundary condition *CX* by comparing two instances which differ only in *CX* and agree in all other circumstances *C*. If in one instance, both *CX* and *A* are present and in the other both *CX* and *A* are absent, then *CX* is causally relevant to *A*. There is a twin method to the method of difference that one might call the *strict method of agreement*, which establishes causal irrelevance, if the change in *CX* has no influence on *A*. Eliminative induction can deal with functional dependencies and an extension of the approach to statistical relationships is straightforward.⁶

Thus, causal (ir-)relevance is a three-place relation: a boundary condition *C* is (ir-)relevant to a phenomenon *A* with respect to a certain background *B* of further conditions that remain constant if causally relevant or are allowed to vary if causally irrelevant. The restriction to a context *B* is necessary because there is no guarantee that in a different context *B**, the causal relation between *C* and *A* will continue to hold. Causal laws established by eliminative induction thus have a distinctive contextual or *ceteris-paribus* character. Extensive information about all potentially relevant boundary conditions in as many different situations as possible is necessary to establish reliable causal knowledge by means of eliminative induction. Exactly this kind of information is provided in data-intensive science.

Eliminative induction corresponds to a difference-making account of causality, which is closely related to the counterfactual approach. However, the truth-value of counterfactuals is now determined via the method of difference or the direct method of agreement, and thus by comparison with actual situations that differ from the counterfactual statement only in terms of irrelevant circumstances, and not by a possible-world semantics as in traditional counterfactual approaches like that of David Lewis.

Obviously, classificatory trees rely on eliminative induction. Thus, to assess their quality, one has to look at the premises required for eliminative methods to yield the correct causes. Partial analyses of this problem are given for example in Keynes (1921, Ch. 22), von Wright (1951, Ch. V), Baumgartner & Grasshoff (2004, Sec. IX 2.4), Pietsch (2014, Sec. 3f). We will again follow the exposition in the last reference. There are at least three main assumptions: (i) determinism, i.e. that the phenomenon *A* is fully determined by boundary conditions *C* and background *B*; (ii) constancy of the background, i.e. that no relevant parameters in the background change when two instances are compared via the method of difference or the strict method of agreement; and finally (iii) an adequate vocabulary, that the parameters *C* reflect suitable causal categories for the given context *B*. Applied to classificatory trees, we can for example say: if there is a single sufficient condition *CX* among the *C* and there is sufficient data in terms of instances of the system in various configurations to avoid spurious correlations, then the classificatory tree algorithm will return *CX* as cause. Certainly, these

⁶ For further discussion, see Pietsch (2014).

assumptions are quite strong. And there are supposedly weaker constraints for causal relations of statistical nature, but this issue goes beyond the scope of the present paper.

We can now identify the elements of theory that have to be presupposed. In particular: (a) one has to know all parameters C that are potentially relevant for the phenomenon A in a given context determined by the background B ; (b) one has to assume that for all collected instances and observations the relevant background conditions remain the same, i.e. a stable context B ; (c) one has to have good reasons to expect that the parameters C are formulated in stable causal categories that are adequate for a specific research question; (d) there must be a sufficient number of instances to cover all potentially relevant configurations of the phenomenon. If such theoretical knowledge can be established, then there is enough data to avoid spurious correlations and to map the causal structure of the phenomenon without further internal theoretical assumptions about the phenomenon.

This motivates and explains the definition of data-intensive science given in Section 2. In particular, premise (I) is the fundamental condition allowing for a strongly inductive approach based on parameter variation. This viewpoint is further corroborated by the fact that in many cases data-driven approaches become effective rather suddenly—a transition point that could be called a *data threshold* (Halevy et al. 2009). Halevy et al. give a plausible explanation for its existence: ‘For many tasks, once we have a billion or so examples, we essentially have a closed set that represents (or at least approximates) what we need, without generative rules.’ (2009, 9) At this threshold, the data represents a large fraction of the relevant configurations of the considered phenomenon.

Of course, in scientific practice full theoretical knowledge a) to d) is rarely available. However, in general, including more potentially relevant parameters C will increase the probability that the actual cause of A might be among them, while admittedly also increasing the probability for spurious correlations, i.e. that boundary conditions accidentally produce the right classification. However, more data in terms of instances of different configurations can reduce the probability for such spurious correlations. Thus, more data in terms of parameters and instances will generally increase the chance that correct causal relations are identified by data-intensive algorithms.

5. Second case study: non-parametric regression

A recent paradigm shift in statistics closely mirrors the change from a hypothesis-directed to a more inductive, data-driven approach. It has been described as a transition from *parametric* to *non-parametric* modeling (e.g. Wassermann 2006; Russell & Norvig 2010, Ch. 18.8), from data to algorithmic models (Breiman 2001), or from model-based to model-free approaches. Since the shift concerns methodology and not theoretical or empirical content, it differs in important ways from scientific revolutions. Nevertheless, the statistics community has experienced over the past two decades some of the social ramifications and ‘culture clashes’ that are typical for scientific paradigm shifts as documented for example in Breiman (2001) or in Norvig’s dispute with Noam Chomsky on data-driven machine translation (Norvig 2011).

This paradigm shift has the following basic features: i) Parametric methods usually presuppose considerable modeling assumptions. In particular, they summarize the data in terms of a ‘small’ number of model parameters specifying for example a Gaussian distribution or linear dependence, hence the name. By contrast, non-parametric modeling presupposes *few modeling assumptions*, e.g. allows for a wide range of functional dependencies or of distribution functions. ii) In non-parametric modeling, predictions are calculated *on the basis of ‘all’ data*. There is no detour over a parametric model that summarizes the data in terms of a few parameters. iii) While this renders non-parametric modeling quite *flexible* with the ability to quickly react to unexpected data, it also becomes extremely *data- and calculation-intensive*. This aspect accounts for the fact that non-parametric modeling is a relatively recent development in scientific method strongly dependent on advances in information technology. It has largely emerged in parallel with the rise of data-intensive science.

Let me give a simple example as an illustration, the comparison between parametric and non-parametric regression. In a parametric univariate linear regression problem, one has reasonable grounds to suspect that a number of given data points $(x_i; y_i)$ can be summarized in terms of a linear dependency: $y = ax + b$. Thus, two parameters need to be determined, offset b and slope a , which are usually chosen such that the sum of the squared deviations $\sum_{i=1}^n (y_i - (ax_i + b))^2$ is minimized.

In non-parametric regression, the data is not summarized in terms of a small number of parameters a and b , but rather all data is kept and used for predictions (e.g. Russell & Norvig 2009, Ch. 18.8.4). A simple non-parametric procedure is *connect-the-dots*. Somewhat more sophisticated is locally weighted regression, in which a regression problem has to be solved for every query point x_q . The y_q -value is determined as $y_q = a_q x_q + b_q$ with the two parameters fixed by minimizing $\sum_{i=1}^n K(d(x_q, x_i))(y_i - (a_q x_i + b_q))^2$. Here, K denotes a so-called kernel function that specifies the weight of the different x_i depending on the distance to the query point x_q in terms of a distance function $d()$. Of course, an x_i should be given more weight the closer it is to the query point.

Let us briefly reflect how these regression methods illustrate the differences between parametric and non-parametric modeling i) to iii). While in parametric regression, linear dependency is presupposed as a modeling assumption, the non-parametric method can adapt to arbitrary dependencies. In parametric regression, the nature of the functional relationship has to be independently justified by the theoretical context, which prevents an automation of the modeling process. Certainly, non-parametric regression also makes modeling assumptions, e.g. a suitable kernel function must be chosen that avoids both over- and underfitting. However, within reasonable bounds the kernel function can be chosen by cross-validation. Since often, predictions turn out relatively stable with respect to different choices of kernel functions, an automation of non-parametric modeling remains feasible.

While non-parametric regression is more flexible than parametric regression, it is also much more data-intensive and requires more calculation power. Notably, in the parametric case, a regression problem must be solved only once. Then all predictions can be calculated from the resulting parametric model. In the non-parametric case, a regression problem must be solved for every query point. In principle, each prediction takes recourse to all the data. While the

parametric model consists in a relatively simple mathematical equation, the non-parametric model consists in all the data and an algorithmic procedure for making predictions.

The main difference in terms of theoretical assumptions is that in parametric regression the type of functional dependency is presupposed in contrast to non-parametric regression. The latter again relies on eliminative induction. Essentially, it constitutes a case of Mill's method of concomitant variations, which derives its inferential power from the method of difference as argued for example in Skyrms (2000, Sec. V.9) and Pietsch (2014, Sec. 3d). Thus, the conditions for identifying a causal relationship are largely the same as those discussed in the previous section—determinism, constancy of the background, and correct causal language—resulting in the same premises in terms of theoretical assumptions a)-d). In particular, when mapping a functional dependency, all causally relevant conditions in the background must remain constant. And there must be sufficient data points such that the functional dependence can be traced in adequate detail.

6. Conclusion: data-intensive science and exploratory experimentation

We are finally in a position to evaluate the claims concerning a theory-free science. In both case studies, certain elements of theory had to be presupposed in order to yield reliable results in terms of causal structure that in turn can underwrite successful prediction and manipulation. In particular, among the considered parameters must be those that are causally relevant for a phenomenon in a considered context and not too many that are causally irrelevant to avoid spurious correlations. Also, the parameters should reflect adequate causal categories. Finally, the collected instances or observations should cover all configurations that are relevant in the given context.

Because these aspects all concern the framing of the problem, one could speak of *external* theory-ladenness. By contrast, there is another kind of theory-ladenness that is largely absent from data-intensive science. For example, in classificatory trees no hypotheses are made about the causal connections that link the various parameters. Equally, in non-parametric regression, no assumptions are presupposed about the functional dependencies between different quantities. Thus, the essential difference in comparison with a hypothesis-driven approach is that not much is presupposed about the *internal* causal structure of the phenomenon. Rather, this structure is mapped from the data by parameter variation.

How novel is this approach? On closer scrutiny, data-intensive science much resembles the practice of exploratory as distinguished from hypothesis-directed experimentation (Steinle 1997, Burian 1997, Waters 2007; cp. also Vincenti 1993, 291). Exploratory experimentation essentially consists in the very same parameter variation of eliminative induction, where the experimenter tries to map the system of interest in all those states that she considers relevant. It is this common methodological core, which links exploratory experimentation and data-intensive science and speaks against the claim, for example by Krohs (2012), that the latter constitutes a novel experimental approach focusing on data-gathering.

Not surprisingly, the debate concerning theory-ladenness in exploratory experimentation parallels the discussion in the present article. For example, Steinle (2005) suggests a

distinction between different kinds of theory-ladenness. According to this view, exploratory experimentation presupposes theoretical knowledge in terms of classification systems or empirical rules, but not in terms of theories that postulate empirically inaccessible abstract entities (285). Steinle refers to Duhem, Hacking and Cartwright as having drawn similar distinctions between an experimental/phenomenological and a theoretical level in scientific theories. Indeed, the distinction between exploratory and hypothesis-driven experimentation fits well with Hacking's (1983) claim that experiments have a life of their own and Cartwright's (1983) position of entity realism, which postulates a causal level in science that is mostly phenomenological and largely independent of the theoretical level.

Building on Burian and Steinle's work, Kenneth Waters emphasizes a subtle difference between 'theory-directed' and 'theory-informed'. While in exploratory experimentation, background theories are used 'to set up experiments, generate data, and draw conclusions', such experiments 'are not "directed" by the aim to test, develop, or otherwise articulate an existing theory or hypothesis.' (2007, 280) Laura Franklin makes a similar point that exploratory experiments are theory-laden in terms of background knowledge, but not in terms of local theories (2005, 891).

These remarks closely parallel the previous discussion regarding external and internal theory-ladenness. The distinction between a phenomenological and a theoretical level is also helpful for the analysis of data-intensive science, which supposedly concerns the phenomenological level regarding local, causal structure of phenomena, but does not rise to the theoretical level.

An important difference between exploratory experimentation and data-intensive science is that in the former, data is usually of experimental nature, while the latter often deals with observational data. But this is largely irrelevant from the perspective of a difference-making account of causation according to which experimental intervention has only pragmatic advantages over observational data. Another difference concerns the complexity of the phenomena. While mapping the causal structure by parameter variation is as old as science itself, carrying it out in the computer can address phenomena that were previously largely inaccessible to causal analysis. This new handle, which data-intensive science provides, for mapping the causal structure of highly complex phenomena will make all the difference to scientific practice.

Acknowledgments

I am grateful to Mathias Frisch, Sabina Leonelli, and Sylvester Tremmel for very helpful insights and discussions.

References

- Anderson, Chris. 2008. "The End of Theory: The Data Deluge Makes the Scientific Method Obsolete." *WIRED Magazine* 16/07.
- Bacon, Francis. 1620/1994. *Novum Organum*. Chicago, IL: Open Court.

- Baumgartner, Michael & Gerd Graßhoff. 2004. *Kausalität und kausales Schließen*. Norderstedt: Books on Demand.
- Breiman, Leo. 2001. "Statistical Modeling: The Two Cultures." *Statistical Science* 16(3):199-231.
- Burian, Richard. 1997. "Exploratory Experimentation and the Role of Histochemical Techniques in the Work of Jean Brachet, 1938–1952." *History and Philosophy of the Life Sciences* 19:27–45.
- Callebaut, Werner. 2012. "Scientific Perspectivism: A Philosopher of Science's Response to the Challenge of Big Data Biology." *Studies in History and Philosophy of Biological and Biomedical Science* 43(1):69-80.
- Cartwright, Nancy. 1983. *How the Laws of Physics Lie*. Oxford: Oxford University Press.
- Franklin, Laura R. 2005. "Exploratory Experiments." *Philosophy of Science* 72:888-899.
- Gray, Jim. 2007. "Jim Gray on eScience: A Transformed Scientific Method." In Tony Hey, Stewart Tansley & Kristin Tolle (eds.). *The Fourth Paradigm. Data-Intensive Scientific Discovery*. Redmond, WA: Microsoft Research.
- Hacking, Ian. 1983. *Representing and Intervening*. Cambridge: Cambridge University Press.
- Halevy, Alon, Peter Norvig & Fernando Pereira. 2009. "The Unreasonable Effectiveness of Data." *IEEE Intelligent Systems* 24(2):8-12.
- Krohs, Ulrich. 2012. "Convenience Experimentation." *Studies in History and Philosophy of Biological and Biomedical Sciences* 43(1), 52–57.
- Laney, Doug. 2001. "3D Data Management: Controlling Data Volume, Velocity, and Variety." Research Report. <http://blogs.gartner.com/doug-laney/files/2012/01/ad949-3D-Data-Management-Controlling-Data-Volume-Velocity-and-Variety.pdf>
- Leonelli, Sabina. 2012. "Making sense of data-driven research in the biological and biomedical sciences." *Studies in History and Philosophy of Biological and Biomedical Sciences* 43, 1–3.
- Mackie, John L. 1980. *The Cement of the Universe*. Oxford: Oxford University Press.
- Mill, John S. 1886. *System of Logic*. London: Longmans, Green & Co.
- Norvig, Peter. 2011. "On Chomsky and the Two Cultures of Statistical Learning." <http://norvig.com/chomsky.html>
- Norvig, Peter. 2009. "All we want are the facts, ma'am." <http://norvig.com/fact-check.html>
- Pietsch, Wolfgang. 2014. "The Nature of Causal Evidence Based on Eliminative Induction." In P. Illari and F. Russo (eds.), *Topoi*. Doi:10.1007/s11245-013-9190-y
- Russell, Stuart & Peter Norvig. 2009. *Artificial Intelligence*. Upper Saddle River, NJ: Pearson.
- Skyrms, Brian. 2000. *Choice and Chance*. Belmont, CA: Wadsworth.
- Steinle, Friedrich. 1997. "Entering New Fields: Exploratory Uses of Experimentation." *Philosophy of Science* 64:S65–S74.
- Steinle, Friedrich. 2005. *Explorative Experimente*. Stuttgart: Franz Steiner Verlag.
- Vincenti, Walter. 1993. *What Engineers Know and How They Know It*. Baltimore: Johns Hopkins University Press.
- Wasserman, Larry. 2006. *All of Nonparametric Statistics*. New York: Springer.
- Waters, C. Kenneth. 2007. "The Nature and Context of Exploratory Experimentation." *History and Philosophy of the Life Sciences* 29(3): 275-284.

von Wright, Georg H. 1951. *A Treatise on Induction and Probability*. New York, NY:
Routledge.