COGENCY Vol. 2, No. 1 (111-121), Winter 2010 ISSN 0718-8285

Gilbert Harman and Sanjeev Kulkarni, *Reliable Reasoning: Induction and Statistical Learning Theory*. A Bradford Book: The MIT Press, 2007, 108 pp., $32.00 (hb), ISBN 978-0-262-08360-7.

Robert Boyd
Department of Philosophy, Fresno City College, Fresno, California, United States.
robert.boyd@fresnocitycollege.edu

Received: 14-03-2010. **Accepted**: 15-06-2010.

Introduction

For those seeking a theoretical discussion regarding inductive reasoning as a reliable tool for drawing inferences, *Reliable Reasoning: Induction and Statistical Learning Theory* by Gilbert Harman and Sanjeev Kulkarni will prove both stimulating and challenging. (An earlier draft of this book was presented as the 2005 Jean Nicod Lectures in Paris by Harman.) The book grew out of an introductory course titled "Learning Theory and Epistemology," which was jointly taught by the Electrical Engineering and Philosophy Departments at Princeton University. The course is described as a low-level undergraduate course serving "as an introduction to aspects of philosophy, computer science, engineering, statistics, and cognitive sciences." (p. ix) This book builds on the work of Vladimir Vapnik and Alexey Chervonenkis.

After a brief overview of the book, I will identify what I believe to be some of the strengths and weaknesses of the book. It should be noted that my evaluation comes from the point of view of one whose primary responsibility has been to teach and develop critical reasoning skills among students. (If a reader of this review is seeking reaction to Harman and Kulkarni's work from other theoreticians working in related fields, I recommend *Abstracta: Linguagem, Mente & Ação*, Special Issue III 2009 (http://www.abstracta.pro. br/english/), which is entirely devoted to their work.)

I believe Harman and Kulkarni's work does contribute to the study of inductive reasoning. Unfortunately, this book has distractions that may cause some readers to fail to recognize its contribution. Some of these distractions will be identified. Finally, one means of evaluating a theoretical work is to consider its viability. Is the position proposed plausible? Does the position promote practical application? To both of these questions, I believe this book is very viable in spite of its major distractions, and I will attempt to illustrate a practical application of their insights.

1. The chapters

In the opening chapter ("The Problem of Induction") Harman and Kulkarni jump directly into the philosophical problem of induction. With little background regarding the problem surrounding inductive reasoning, the authors state that the problem of induction is "about the *reliability* of inductive inferences" (p. 1). As they develop the problem of induction, they are unwilling (and rightly so) to suggest that the problem is simply that induction is not deductive. "Deductive logic is a theory of what follows from what, not a theory of reasoning" (p. 6). Their concern focuses upon a process of reasoning, which is induction. At the end of the opening chapter they link this process of reasoning with the question of reliability, and they propose that statistical learning theory may provide a method for identifying the reliability of induction.

Chapter two ("Induction and VC Dimension") begins with a critical distinction. A distinction must be made, they claim, between rules of classification and inductive methods for finding such rules. "Rules of classification ... are rules for using observed *features* of items to classify them or to estimate the values of a real variable. Inductive methods for finding such rules are methods for using *data* to select such rules of classification or estimation" (p. 29). This is accomplished by revealing "a pattern in the data that can then be used to classify new cases ..." (p. 30). This second chapter focuses on enumerative induction as a method for finding rules of classification, and it provides a discussion of pattern classification, real variable values, probability distribution as well as several other issues before suggesting how statistical learning theory can assess inductive methods, specifically enumerative induction.

Chapter three ("Induction and "Simplicity"") attempts to illustrate how statistical learning theory might shed light on two philosophical issues: Goodman's "new riddle of induction" and Popper's scientific method of falsificationism. These two issues are discussed as a means to compare enumerative induction with an alternative method such as simplicity that seeks to give some type of ordering over alternative hypotheses. Here simplicity is understood as a "principled way to prefer one hypothesis over another skeptical hypothesis that is empirically equivalent to it" (p. 55).

In the final chapter ("Neural Networks, Support Vector Machines, and Transduction") the authors "briefly sketch some applications of statistical learning theory to machine learning, including perceptrons, feed-forward neural networks, and support vector machines" (p. 77). They also discuss possible application of "transduction." Whereas induction uses labeled data to formulate rules of classification, transduction uses new unlabeled data. "The theory of transduction suggests new models of how people sometimes reason" (p. 98).

2. Strengths

While the field of critical reasoning or informal logic has developed much since its development in the 1960s and 1970s, we still have not fully come to grips with inductive reasoning, how to evaluate it or how to improve its reliability. At the core of this deficit is the lack of a satisfactory "metatheory" for induction. Whereas for deductive logic there are many fine books following the example of Geoffrey Hunter's (1971) *Metalogic: An Introduction to the Metatheory of Standard First Order Logic*, meta works on induction are still needed. (Hunter, himself, expressed a possible "extension of logical theory" as he considered the following propositions: "A man who knows of at least one case of an X being a Y, and who does not know of any positive reason for thinking that an X might not be a Y, has some reason for thinking that all X's are Y's." (Hunter, 1965).)

Given this deficit in the field of inductive reasoning, Harman and Kulkarni provide an important starting point for the development of a metatheory for induction. Their insight of applying statistical learning theory to provide a foundation for identifying the reliability of induction is significant. It should be noted that while they call for the application of statistical learning theory, they fail, in spite of chapter three, to show how this is done in regards to philosophical issues. Their project is to provide an inductive method which will then be used to identify rules of classification or estimate values of real variables. While their book itself is not the needed "meta-inference: an introduction to the metatheory of inductive reasoning," it may lay the foundation like Russell's (1919) *Introduction to Mathematical Philosophy* laid for subsequent metalogics. As such, Harman and Kulkarni have provided an important contribution.

Another strength of this work is its development of the philosophical problem of induction in chapter one. While inductive reasoning has many applications and hence can be used on various problems, the philosophical problem of induction is its reliability. This reliability cannot be related to that of deductive logic. As the authors clearly state, "it is a mistake to describe the problem of inductive reliability by comparison with deductive reliability. Deductive rules are rules about what follows from what; they are not rules about what can be inferred from what" (p. 9). The rules of inference must be about inductive reasoning. For a complete system of "logic" (not Harman and Kulkarni's term), we need both rules of derivation and rules of inference. This distinction as well as how it relates to the problem of induction is valuable.

Just as valuable is their discussion of pattern classification. As they tell us, "[a]n inductive method is a principle for finding a *pattern* in the data that can then be used to classify new cases or to estimate the values of a real variable" (p. 30). While the understanding that the inductive method seeks patterns is not new, they suggest non-qualitative means of representing those patterns. For example, it is now common to teach qualitative methods for portraying patterns seen in a piece of inductive reasoning. This method often entails a written narrative, a telling of a story. Harman and Kulkarni's discussion of pattern classification points to quantitative representations of the data. They show this quantitative representation in the form of graphs with x and y axes, but this process opens the door to graphing inductive reasoning into other mathematical methods of presentation such as fractals. Given this extension, we might be able to eventually visually tell the difference between an inductively strong reasoning and an inductively weak reasoning.

3. Distractions

While this book does have the above strengths and does make a positive contribution to the study of inductive reasoning, it does present a number of major distractions, which must be acknowledged. The first thing that confronts its reader is that the authors seem to have misidentified their audience. In their introduction they set forth the context in which this book developed. This context would suggest that their audience is lower-division undergraduate students. While it was written as an introduction to the subject, it assumes its readers are well versed with each of the topics discussed: problem of induction; statistical learning theory; Goodman and Popper's philosophy, etc.

In order for this text to be successfully used as a textbook for lower-division students with no particular prerequisites "other than some analytical skills and intellectual curiosity" (p. ix), an instructor must heavily supplement it with required background knowledge and skills. This, of course, can be done, and I am sure is being done in the class at Princeton where Harman and Kulkarni teach. If this text is not primarily aimed at lower-division students, then who is its intended audience? Given the type of assumptions the authors make about their readers, e.g., basic knowledge of philosophical issues and statistical probability skills, it appears they are primarily writing for PhD level students or for other scholars who are already working in the area of the process of reasoning and its reliability. This is not to suggest the book has no value as an undergraduate textbook; rather, because of critical assumptions made by the authors regarding their audience, they run the risk of minimizing the impact of their insights.

Like many other very influential works in the field of reasoning, such as Toulmin's (1958) *The Use of Arguments* or Plantinga's (1974) *The Nature of Necessity*, *Reliable Reasoning: Induction and Statistical Learning Theory* is written as a theoretical proposal. It offers those of us working in the various fields related to reasoning a theoretical framework in which we can construct the various problems we are currently focusing. In this sense, Harman and Kulkarni have provided a very valuable service. However, I believe Harman and Kulkarni could have strengthened their proposal by providing better and more complete examples of their theory in application. As I understand their project, they believe that statistical learning theory, when coupled with enumerative induction, can provide a method for determining rules of induction that are reliable. Even when they finally offer illustrations of application, their illustrations are underdeveloped and leave the reader without direction.

A third major distraction within *Reliable Reasoning: Induction and Statistical Learning Theory* is the sentence structure and language used throughout the text. There are terms used that one cannot assume a lower-division undergraduate would be familiar with. For example, Harman and Kulkarni refer to Bayes Rule throughout the book, but at no place is this rule presented. Moreover, the VC Dimension (named after Vapnik and Chervonenkis 1971) is crucial in key passages of the text, but it remains undefined. Finally, Harman and Kulkarni rightly claim that the philosophical problem of induction focuses on the issue of reliability. The notion of reliability is even part of the title of this book, yet the authors fail to define what they mean by reliable or exactly what constitutes reliable reasoning. These are only three examples of when the language chosen becomes an obstacle rather than a gateway to greater insights for the reader. Even if we assume the text was actually written not for lower-division undergraduates, but for scholars working in the field, the sentence structure frequently found in the text is problematic at best. Consider the following sentences:

> Vapnik and Chervonenkis show that the method of empirical risk minimization, when used to select rules of classification, has the following property. If, and only if, the VC dimension of C is finite, then no matter what the background probability distribution, as more and more data are obtained, with probability approaching 1, enumerative induction leads to the acceptance of rules whose expected error approaches the minimum expected error for rules in C. (p. 56)

While the first sentence is unproblematic and is presented here only to provide some context for the second sentence, the second sentence is rather difficult to parse. If this were the only instance of problematic sentence structure, then I could be charged with being petty, but throughout the book, and at critical points, such statements are provided. Of course, this is not a problem for just Harman and Kulkarni. While there are many fine examples of well-written analytical philosophy, current standards in the tradition do not always contribute to clarity. As scholars, in whatever field we work, we all would do well to occasionally reread Russell's (1954) essay "How I Write."

4. Evaluation

I suggested that one way of evaluating a theoretical work is to ask whether it is viable in the sense that it promotes an extension of itself. In the following I present two passages; each are examples of inductive reasoning that may be encountered in every day type of reasoning. (The first passage is intuitively strong, while the second is obviously fallacious.) Along with each passage I have provided a possible assessment of the reasoning. This assessment will first put the reasoning in standard form. From this we can identify the type of enumerative induction and the relevant criteria for that type. (Note that I will not provide here a justification for this labeling nor for the selected criteria; that is outside the scope of this review, but common in informal logic texts, e.g., Boyd (2003) or Salmon (1989). Enumerative induction refers to a class of reasoning and is frequently sub-divided into specific types, i.e., simple enumerations, inductive generalizations, arguments from analogy, statistical syllogisms, etc.) Following this I will present a very brief qualitative assessment. Finally, I will suggest a quantitative assessment based upon the qualitative evaluation. (The quantitative assessments are assigned based upon key hedging or frequency terms used in the qualitative evaluations and their frequency estimates. Whereas McNeill and Freiberger (1993) present median results from studies by Ray Simpson (1944) and Milton Hakel (1968), I have assigned frequency ranges. Clearly, more work, building upon that of Lakoff (1973), must be done regarding the use of hedging terms for this type of reasoning assessment to be successful.) The point of this presentation is not to provide a complete and proper assessment, but only to illustrate the possibility of moving from some level of qualitative assessment to one that is quantitative, which is motivated by the theoretical work of *Reliable Reasoning: Induction and Statistical Learning Theory.*

Passage 1:

> Board members need policies, training activities, guidelines to ensure

> that if and when sexual harassment occurs, school officials are prepared to deal with it. . . . Furthermore, if a sexual harassment case reaches the courts, the grievance and rectifying procedures provide some legal protection, proving that your board made a good-faith effort to prevent sexual harassment among employees. You also might find your efforts make it possible to prevent it. (Decker, 1989)

Standard form for reasoning

> (R1) If SH policy is in place, then when SH occurs officials are prepared to deal with it.
> (R2) If SH case reaches courts and a SH policy is in place and good-faith efforts have been made to prevent SH, then the courts tend to provide some legal protection for such schools.
> (R3) If SH policy is in place, it may prevent SH.
> Inference drawn: all school boards should have a SH policy.

(The phrase 'Inference drawn' is used following Harman and Kulkarni's position that the term 'inference' should not be linked with deductive logic, but exclusively to inductive reasoning. (pp. 5-9) Furthermore, they claim that one should not refer to inductive arguments, but to inductive reasoning. (p. 7) As a result, I avoid claiming the inference to be a 'conclusion', which is the typical way we refer to the claim statement being drawn from evidence.)

This is an Inductive Generalization, and the relevant criteria are Sample Diversity, Sample Size, and Other considerations.

Qualitative analysis: The reasoning illustrates a significant level of diversity since each of the three pieces of data (i.e., R1 to 3) represents different relevant arenas of interest to school boards considering sexual harassment. Because of this level of diversity, we believe that sample diversity in this passage is *very often* satisfied. Given the qualitative analysis, we might suggest that the quantitative probability of diversity being adequately satisfied falls within a range of 88-93%.

When considering the sample size, we note that the reasoning is based on only three pieces of data; however, sample size, while normally focusing on quantitative weight, does have a qualitative nature. Since each piece of data on its own in the given context would support the inference drawn, we suggest that sample size in this passage could be understood as *almost always* being satisfied. Quantitatively we suggest a range of 95-98% probability that sample size is satisfied.

Regarding other considerations, we notice that the inference – while embracing "all" school boards – is proposing a suggestion, i.e., "should," which is a much weaker proposal than a stronger proposal of "must." Since the inference is a weaker form, there is a greater probability that it is supported by the data. We conclude that this third criterion, given its weakened form, is *usually* satisfied. Quantitatively, we assign this criterion an 83-85% probability.

Given this analysis, we can calculate the overall probability of the inference being supported by the reasons as: lower range - .88 x .95 x .83 = 69% and upper range - .93 x .98 x .85 = 77% probability. So this reasoning illustrates a range between 69% and 77% probability of inferring that all school boards need sexual harassment policies.

Passage 2:

> "Smoking by pregnant women may result in fetal injuries, premature birth, and low birth weight." (Surgeon General's warning on a pack of cigarettes.) Susan is a 95 year old smoker. Probably, her next child will be born prematurely.

Standard form for reasoning

> (R1) Smoking by pregnant women may result in fetal injuries, ...
> <u>(R2) Susan is a 95 year old smoker.</u>
> Inference drawn: probably Susan's next child will be born prematurely.

This is a Simple Enumeration, and the relevant criteria are Sample Size, Total Evidence, and Other considerations.

Qualitative analysis: We can assume that the tobacco industry was forced to place this warning on its product due to an extremely large amount of evidence that supported the warning. Sample size is probably *very often* satisfied in such cases. The quantitative range of 88-93% should be assigned for the probability of sample size being satisfied.

However, total evidence is extremely problematic. The evidence regarding the probability of Susan becoming pregnant given her age of 95 suggests this criterion has been violated. Since we are dealing with Susan and not the Sarah of biblical legend, we can conclude that in this situation total evidence would *almost never* be satisfied. Thus, total evidence here comprises all relevant reasons. We suggest there is only a 2-3% chance of total evidence being supported in this inference.

Regarding other considerations, because of the qualifier "probably" in the inference drawn, the criterion ranks higher than it would in a stronger form of the inference without the qualifier. In this case, we believe the criterion will *usually* be satisfied. We suggest an 83-85% range for this criterion.

Given this analysis, we can calculate the overall probability of the inference following from the data as: lower range - .88 x .02 x .83 = 1% and upper range - .93 x .03 x .85 = 2% probability. So this reasoning illustrates a range between 1% and 2% probability of supporting the inference that Susan's next child will be born premature. This low range points to the fallacy of incomplete evidence committed in the reasoning.

These two problems illustrate how using Harman and Kulkarni's lead we might move beyond a mere qualitative analysis of inductive reasoning to a quantitative analysis which can be converted to and portrayed in graph form or fractal for a visualization of the reasoning. The first problem could visually portray a strong to moderate inductive inference, whereas the second would portray a fallacious inductive inference. While the analysis provided for these two problems should not be taken seriously, they do illustrate that *Reliable Reasoning: Induction and Statistical Learning Theory* provides a valuable service as a theoretical work that promotes further research in the field of inductive reasoning.

Works Cited

Abath, André and Leonardo de Mello Ribeiro (eds.). *Abstracta: Linguagem, Mente e Ação.* Special Issue III, 2009 (http://www.abstracta.pro.br/ english/).

Boyd, Robert. *Critical Reasoning and Logic* (pp. 104-148). Upper Saddle River, NJ: Prentice Hall, 2003.

Decker, Robert. "Can schools eliminate sexual harassment?" *The Educational Digest*, January 1989: 59-62.

Hakel, Milton. "How often is often?" *American Psychologist* 23 (1968): 533-534.

Hunter, Geoffrey. "A possible extension of logical theory?" *Philosophical Studies* XVI (6) (1965): 81-88.

Hunter, Geoffrey. *Metalogic: An Introduction to the Metatheory of Standard First Order Logic.* Berkeley: University of California Press, 1971.

Lakoff, George. "Hedges: A study in meaning criteria and the logic of fuzzy concepts". *Journal of Philosophical Logic* 2 (1973): 458-508.

McNeill, Daniel and Paul Freiberger. *Fuzzy Logic* (pp. 67-68). New York: A Touchstone Book: Simon and Schuster, 1993.

Plantinga, Alvin. *The Nature of Necessity.* New York: Clarendon Press, Oxford, 1974.

Russell, Bertrand. "How I Write", *The Writer.* September 1954. (http://www.solstice.us/russell/write.html).

Russell, Bertrand. *Introduction to Mathematical Philosophy.* New York: A Touchstone Book: Simon and Schuster. (First published by George Allen and Unwin Ltd., 1919.)

Salmon, Merrilee. *Introduction to Logic and Critical Thinking,* 2nd edition (pp. 64-106). San Diego: Harcourt Brace Jovanovich, 1989.

Simpson, Ray. "The specific meanings of certain terms indicating differing degrees of frequency". *The Quarterly Journal of Speech* 30 (1944): 328-330.

Toulmin, Stephen. *The Uses of Argument.* Cambridge: Cambridge University Press, 1958.

Vapnik, Vladimir and Alexey Chervonenkis. "On the uniform convergence of relative frequencies of events to their probabilities". *Theory of Probability and its Applications* 16 (2) (1971): 264-280.