Smart Inductive Generalizations are Abductions

(Parts are adapted from Josephson & Josephson, 1994, Chapter 1,
"Conceptual analysis of abduction." Used with permission.)

John R. Josephson (jj@cis.ohio-state.edu)
The Ohio State University
Laboratory for AI Research
Computer and Information Science Department
Columbus, Ohio, USA

Abstract

This paper describes abduction as 'inference to the best explanation' and argues that "smart" inductive generalizations are a special case of abductions. Along the way it argues that some good explanations are not proofs and some proofs are not explanations, concluding that explanations are not deductive proofs in any particularly interesting sense. An attractive alternative is that explanations are assignments of causal responsibility. Smart inductive generalizations can then be seen to be abductions wherein the frequency in a statistical sample is best explained by a frequency in a parent population along with the method of drawing the sample.

A distinctive pattern of inference

To postpone entanglements with the abundant confusions surrounding various uses of the term "abduction," for which Peirce himself seems largely to be responsible, and to proceed as directly as possible to engage the basic logical and computational issues, let us begin by examining a pattern of inference I will call "inference to the best explanation" and abbreviate as "IBE" ("IBEs" for the plural).

IBEs follow a pattern something like this:

D is a collection of data (facts, observations, givens),
H explains D (would, if true, explain D),
No other hypothesis explains D as well as H does.

Therefore, H is probably correct.

The strength the conclusion depends on several factors, including:

* how good H is by itself, independently of considering the alternatives,
* how decisively H surpasses the alternatives, and
* how thorough the search was for alternative explanations.

At the outset, I will not suggest that anyone actually reasons this way. Rather, I suggest that it is a *good way* to reason, that reasoning this way has epistemic merit, that reasoning this way yields conclusions that arrive with some evidential force.

In fact, we can observe that people quite commonly justify their conclusions by direct or barely disguised appeal to IBE. It is already commonly believed to be a good way to reason. IBE justifications would not work to persuade, if the listener could not already be expected to be susceptible to the appeal. Why people are persuaded by IBE justifications is a matter for interesting speculation - perhaps it is somehow built into the human mind, and perhaps this is good design - but *that* they are persuaded by IBE is, I think, a matter of empirical fact. It seems that people intuitively recognize many of the factors, such as those I have mentioned, that govern the strength of the conclusions of IBEs. Beyond that, people sometimes actually come up to the standards set by IBE in their actual reasoning, and then they can be reasonably said to be “behaving intelligently” (in that respect). So we may remove our initial reservation about whether people actually reason in accordance with IBE, and agree that they commonly do so; it is part of being intelligent, part of being “smart.” I refer the reader to Harman (1965), Lipton (1991), and Josephson & Josephson (1996) for more extensive discussions of the epistemic virtues of this pattern of inference.

I do not suggest that the description of the IBE pattern that I have given here is perfect, or precise, or complete, or the best possible description of it. I suggest only that it is close enough so that we can recognize it as distinctive, epistemically forceful, ubiquitous, and smart.

Consider another pattern of inference, which I will call “backward modus ponens,” that has a pattern as follows:

p --> q
q

Therefore, p.

The arrow, "-->", here may be variously interpreted, so let us just suppose it to have more or less the same meaning as the arrow used in schematizing:

p --> q
p

Therefore, q.

This one is modus ponens, and this one is smart. Modus ponens has some kind of intuitively visible epistemic force. In contrast, backward modes ponens is obviously fallacious. (Copi calls it “the fallacy of affirming the consequent.” Copi [reference]). By itself, backward modus ponens is not smart, although reasoning in accordance with it may be smart for other reasons, and there may be special contexts in which applying backward modus ponens is smart.

It has become common in AI to identify “abduction” with backward modus ponens, or with backward modus ponens together with constraints on the set of conclusions that are allowable, or with backward modes ponens with syntactic constraints. There is a burden on those who study restricted forms of backward modus ponens to show us the virtues of their particular forms -- they need to show us how they are smart. I suggest that they will find that *backward modus ponens is smart to the degree that it approximates, or when it* *is controlled and constrained to approximate, or when it implements, inference to the best explanation*..

Sometimes a distinction has been made between an initial process of coming up with explanatorily useful hypothesis alternatives and a subsequent process of critical evaluation wherein a decision is made as to which explanation is best. Sometimes the term "abduction" has been restricted to the hypothesis-generation phase. Peirce himself commonly wrote this way, although at other times Peirce clearly used the term "abduction" for something close to what I have here called "inference to the best explanation." (For a discussion of Peirce's views on abduction please see Flach [reference], and for a detailed scholarly examination of Peirce's writings on abduction please see Fann [reference].)

Sometimes "abduction" has been identified with the *creative* generation of explanatory hypotheses, even sometimes with the creative generation of ideas in general. Kruijff suggests that, besides the creativity of hypotheses, the surprisingness of what is to be explained is at the core of abduction's ubiquity and of its relation to reality (Kruijff reference). It is clear that there is much expected utility in trying to explain things that are surprising. Surprise points out just where knowledge is lacking, and when a failed expectation has distinctly pleasant, or more likely, unpleasant, effects, there may well be something of practical important to be learned. But one may also wonder about, and seek explanations for, things that are not ordinarily surprising, and which only become "surprising" when you wonder about them, when you recognize that in some way things could be different. "Why do things fall?" "Why do people get angry?" "Why do arctic foxes have white coats in winter?" None of these are unexpected, all present openings for new knowledge. Clearly, neither novelty of hypothesis nor surprise at the data are essential for an IBE to establish its conclusion with evidential force. "Who forgot to put the cheese away last night?" "Billy again." & etc.

While the creative generation of ideas is certainly virtuous in the right contexts, and useful for being smart, it is necessary for creative hypotheses to have some plausibility, some chance of being true, or some pursuit value, before creativity can make a genuine contribution to working intelligence. Generating low value creative explanatory hypotheses is in itself a dis-virtue in that time, attention, and other cognitive or computational resources must then be expended in rejecting these low value hypotheses so that better hypotheses may be pursued. Too much of the wrong kind of creativity is a drain on intelligence, and so is not smart. Generation of hypotheses, without critical control, is not smart. Generation of hypotheses, as a pattern of inference, in and of itself, is not smart.

Generation of *plausible* explanatory hypotheses, *relevant* to the current explanatory problem, is smart. Yet prescreening hypotheses to remove those that are implausible or irrelevant mixes critical evaluation into the hypothesis-generation process, and so breaks the separation between the process of coming up with explanatory hypotheses and the process of critical evaluation. Furthermore, evaluating one or more explanatory hypotheses may require (according to IBE) that alternative explanations are generated and considered and that a judgment is made concerning the thoroughness of the search for alternative explanations. Again we have a breakdown in the separation of the processes of hypothesis generation from the processes of critical evaluation. Either type of process will sometimes need the other as a subprocess. Moreover, either process might precompile the use of the other process, so that it is not invoked explicitly at run time, but instead is "compiled out" and only implicit. A hypothesis generation mechanism might implicitly use criticism (it must do so if it is to be smart), and criticism might implicitly use hypothesis generation (for example by looking for "pathognomonic" confirming signs, which are already known to have only one possible explanation). Thus, I conclude that it is best not to put much weight on distinguishing hypothesis generation from critical evaluation, and in any case, hypothesis generation by itself is not smart.

From the foregoing discussion it appears that IBE is distinctive, evidentially forceful, ubiquitous, and smart; and that no other proposed definition or description of the term "abduction " has all of these virtues. Thus it seems that IBE is our best candidate as a description of what is at the epistemological and information-processing core of the family of patterns collected around the idea of abduction. I therefore claim the term "abduction" for IBE, and in the remainder of this paper, *by "**abduction**" I mean "inference to the best explanation."*

Some authors characterize abduction as reasoning from effects to causes, a view to which we will return to later in the paper (e.g., Poole [reference]). For now, I would just like to point out that, at least, *abduction is a good way to be effective in reasoning from effects to causes.* If a putative cause is the best explanation, then there is good evidence that the putative cause is the true cause.

Up to now the discussion has mainly focused on abduction as an argument pattern -- as a pattern of evidence and justification -- although we have briefly touched on a process-oriented view of abduction in our discussion of the separability of hypothesis generation and evaluation, and in other hints about what it takes to be smart. Flach has distinguished Peirce's early views on abduction from his later more mature views and has characterized these as "syntactical" and "inferential" views of abduction, the latter being more process oriented (Flach reference). Öztürk has distinguished inference "as an evidential process," which is concerned with the value of conclusions either in security or in productivity, from inference "as a methodological process," which emphasizes the role of inferences in the economy of processes of inquiry or the uses of inferences in support of other tasks (Öztürk reference).

I suggest that it will help clarify matters to distinguish abduction *as an argument pattern*, from abduction *as a reasoning task*, from abduction *as a reasoning process*. An information-processing task sets up a goal to accomplish, which may be described independently of its means of accomplishment, that is, a task may be described separately from the available algorithms, mechanisms, strategies, implementations, and processes that will be needed to accomplish it. (See Lucas [reference] for a method-independent account of diagnosis as a task.) These three perspectives -- justification, task, and process -- are conceptually tightly interconnected. An abductive reasoning task, prototypically, is one that has the goal of producing a satisfactory explanation, which is an explanation that can be confidently accepted. An explanation that can be confidently accepted is one that has strong abductive justification. Thus, a prototypical abductive task aims at setting up strong abductive justifications. Information processing that is undertaken for the purpose of accomplishing a prototypical abductive task, that is, of producing a satisfactory explanation, may reasonably be considered to be an "abductive reasoning process." Thus, from a processing perspective it makes sense to think of abductive reasoning as comprising the whole process of generation, criticism, and acceptance of explanatory hypotheses.

Note that the abductive justifications set up by abductive reasoning might be explicit, as when a diagnosis can be justified, or they might arise implicitly as a result of the functioning of an "abductively effective" mechanism, such as, perhaps, the human visual system, or human language understanding, or an effective neural-net diagnostic system. Note also that the conclusions of abductive arguments (and correspondingly, the accomplishments of abductive tasks, and the results of abductive reasoning processes) may be either general or particular propositions. Sometimes a patient's symptoms are explained; sometimes an empirical generalization is explained by an underlying causal mechanism (e.g., universal gravitation explains the orbits of the planets).

The account of abduction that has been sketched so far in this paper still has two large holes: (1) what is an explanation? and (2) what makes one explanation better than another?

I will not attempt to fill the second hole in this paper -- the literature on the subject is vast. I will simply mention some desirable features of explanatory hypotheses: consistency, plausibility, simplicity, explanatory power, predictive power, precision, and theoretical promise. To begin to fill the first hole, let us ask: what conception of explanation is needed for understanding abduction?

Explanations are not proofs

There have been two main traditional attempts to analyze explanations as deductive proofs. By most accounts, neither attempt has been particularly successful. First, Aristotle maintained that an explanation is a syllogism of a certain form that also satisfies various informal conditions, one of which is that the middle term of the syllogism is the cause of the thing being explained (Aristotle reference). More recently (considerably) Hempel (1965) modernized the logic and proposed the “covering law” or “deductive nomological” (D-N) model of explanation. (For a brief summary of deductive and other models of explanation see Bhaskar (1981); for a history of recent philosophical accounts of explanation, see Salmon (1990); and for a brief discussion of the inadequacies of the D-N model see Psillos (reference).) The main difficulty with these accounts (besides Hempel’s confounding the question of what makes an ideally good explanation with the question of what it is to explain at all) is that being a deductive proof is neither necessary nor sufficient for being an explanation. Consider the following:

QUESTION: Why does he have burns on his hand?

EXPLANATION: He sneezed while cooking pasta and upset the pot.

The point of this example is that an explanation is given, but no deductive proof, and although it could be turned into a deductive proof by including additional propositions, this would amount to gratuitously completing what is on the face of it an incomplete explanation. Real, non-ideal explanations are almost always incomplete. Under the circumstances (incompletely specified) sneezing and upsetting the pot were presumably *causally sufficient* for the effect, but this is quite different from being *deductively sufficient*. For another example, consider that the flu hypothesis explains the body aches, but often people have flu without body aches, so having flu does not imply having body aches. The lesson is that an explanatory hypothesis need not deductively entail what it explains.

The case that explanations are not necessarily deductive proofs is made even stronger when we consider psychological explanations, where there is presumptively an element of free will, and explanations that are fundamentally statistical, where, for example, quantum phenomena are involved. In these cases it is clear that causal determinism cannot be assumed, so the antecedent conditions, even all antecedent conditions together, known and unknown, cannot be assumed to be causally sufficient for the effects.

Conversely, many deductive proofs fail to be explanations of anything. For example classical mechanics is deterministic and time reversible, so an earlier state of a system can be deduced from a later state, but the earlier state cannot be said to be explained thereby. Also, *q* can be deduced from '*p* and *q*' but is not thereby explained. Many mathematicians will at least privately acknowledge that some proofs establish their conclusion without giving much insight into *why* the conclusions is true, while other proofs give richer understanding. So it seems that, even in pure mathematics, some proofs are more explanatory than others.

We are forced to conclude that explanations are not deductive proofs in any particularly interesting sense. Although they can often be presented in the form of deductive proofs, doing so does not succeed in capturing anything essential or especially useful. Thus the search for a proof of D is not the same as the search for an explanation of D, instead it is only a traditional, but seriously flawed, approximation of it.

Explanations give causes

An alternative view is that an explanation is an assignment of causal responsibility; it tells a causal story. Finding possible explanations is finding possible causes of the thing to be explained. It follows that abduction, as a process of reasoning to an explanation, is a process of reasoning from effect to cause. (Ideas of causality and explanation have been intimately linked for a very long time. For a well-developed historical account of the connections, see Wallace (1972, 1974).)

It appears that "cause" for abduction must be understood somewhat more broadly than its usual senses of mechanical, or efficient, or event-event causation. To get some idea of a more expanded view of causation, consider the four kinds of causes according to Aristotle: efficient cause, material cause, final cause, and formal cause (Aristotle, *Physics*, bk. 2, chap. 3). Consider the example of my coffee mug. The *efficient cause* is the process by which the mug was manufactured and helps explain such things as why there are ripples on the surface of the bottom. The *material cause* is the ceramic and glaze, which compose the mug and cause it to have certain gross properties such as hardness. The *final cause* is the end or purpose, in this case to serve as a container for liquids and as a means of conveyance for drinking. A final-cause explanation is needed to explain the presence and shape of the handle. *Formal cause* is somewhat more mysterious -- Aristotle is hard to interpret here -- but it is perhaps something like the mathematical properties of the shape, which impose constraints resulting in certain specific other properties. That the cross-section of the mug, viewed from above, is approximately a circle, explains why the length and width of the cross-section are approximately equal. Apparently, the causal story told by an abductive explanation might rely on any type of causation. What the types of causation and causal explanation are, remains unsettled, despite Aristotle's best efforts and those of many other thinkers. The point here is that a narrow view of causation makes progress harder by obscuring the degree to which various forms of causal reasoning are fundamentally similar.

When we conclude that data D is explained by hypothesis H, we say more than just that H is a cause of D in the case at hand. We conclude that among all the vast causal ancestry of D we will assign responsibility to H. Commonly, our reasons for focusing on

H are pragmatic and connected rather directly with goals of producing or preventing *D*. We blame the heart attack on the blood clot in the coronary artery or on the high-fat diet, depending on our interests. Perhaps we should explain the patient's death by pointing out that the patient was born, so what else can you expect but eventual death? We can blame the disease on the invading organism, on the weakened immune system that permitted the invasion, or on the wound that provided the route of entry into the body. I suggest that it comes down to this: the things that will satisfy us as accounting for *D* will depend on why we are trying to account for *D*; but the only things that count as candidates are parts of what we take to be the causal ancestry of *D*.

Explanations give causes. Explaining something, whether that something is particular or general, gives something else upon which the first thing depends for its existence, or for being the way that it is. The bomb explosion explains the plane crash. The mechanisms that connect the ingestion of cigarette smoke with effects on the arteries of the heart, explain the statistical association between smoking and heart disease. It is common in science for an empirical generalization, an observed regularity, to be explained by reference to underlying structure and mechanisms. Explainer and explained, explanans and explanandum, may be general or particular. Accordingly, abductions may apply to, or arrive at, propositions that either general or particular. Computational models of abduction that do not allow for this are not fully general.

As I have argued, explanations are not deductive proofs in any particularly interesting sense. Although they can often be presented in the form of deductive proofs, doing so seems not to capturing anything essential, or especially useful. Thinking of explanations as proofs tends to confuse causation with logical implication. To put it simply: causation is in the world, implication is in the mind. Of course, mental causation exists (e.g., where decisions cause other decisions), which complicates the simple distinction by including mental processes in the causal world, but that complication should not be allowed to obscure the basic point, which is not to confuse an entailment relationship with the objective, causal grounds for that relationship. Deductive models of causation are at their best when modeling deterministic closed-world causation, but this is too narrow for most purposes. Moreover, even there it is dangerous, since one must be careful to exclude non-causal and anti-causal (effect-to-cause) conditionals from any knowledge base. (Pearl (reference) has pointed out the significant dangers of unconstrained mixing of cause-to-effect with effect-to-cause reasoning.)

Per se, there is no reason to seek an implier of some given fact. The set of possible impliers includes all sorts of riff raff, and there is no obvious contrast set at that level to set up reasoning by exclusion. But there is a reason to seek a possible cause: broadly speaking, because knowledge of causes gives us powers of influence. And the set of all possible causes (at the level we are interested in) does constitute a contrast set for reasoning by exclusion. Common sense takes it on faith that everything has a cause. (Compare this with Leibniz's Principle of Sufficient Reason.) There is no (non-trivial) principle of logic or common sense that says that everything has an implier.

In the search for the interesting causes of our observations we may set up alternative explanations and reason by exclusion. Thus, effect-to-cause reasoning is not itself the same as abduction, rather, effect-to-cause reasoning is what abduction is for.

Smart inductive generalizations are abductions

The work "induction" has had no consistent use, either recently or historically. Sometimes writers have used the term to mean all inferences that are not deductive, sometimes they have specifically meant inductive generalizations, and sometimes they have meant next-case inductions as in the philosophical "problem of induction" as put by David Hume. C. S. Peirce seems to have sometimes used the term to mean the process of testing a hypothesis by generating predictions and evaluating those predictions empirically. An ***inductive generalization*** *is an inference that goes from the characteristics of some observed sample of individuals to a conclusion about the distribution of those characteristics in some larger population.* Examples include generalizations that arrive at categorical propositions (All A's are B's) and generalizations that arrive at statistical propositions (71% of A's are B's, Most A's are B's). A common form of inductive generalization in AI is called "concept learning from examples," which may be supervised or unsupervised. Here the learned concept generalizes the frequencies of occurrence and co-occurrence of certain characteristics in a sample, with the intention to apply them to a larger general population, which includes unobserved as well as observed instances.

I will argue that it is possible to treat every "smart" (i.e., reasonable, valid, strong) inductive generalization as an instance of abduction, and that analyzing inductive generalizations as abductions shows us how to evaluate the strengths of these inferences.

First we note that many possible inductive generalizations are not smart.

> This thumb is mine & this thumb is mine.
> --
> Therefore, all thumbs are mine.
>
> All observed apples are observed.
> --
> Therefore, all apples are observed.

Russell's example: a man falls from a tall building, passes the 75th floor, passes the 74th floor, passes he 73rd floor, is heard to say, "so far, so good."

Harman (1965) pointed out that it is useful to describe inductive generalizations as abductions because it helps to make clear when the inferences are warranted. Consider the following inference:

All observed A 's are B 's

Therefore, All A 's are B 's

This inference is warranted, Harman writes, ". . . whenever the hypothesis that all A 's are B 's is (in the light of all the evidence) a better, simpler, more plausible (and so forth) hypothesis than is the hypothesis, say, that someone is biasing the observed sample in order to make us think that all A's are B's. On the other hand, as soon as the total evidence makes some other competing hypothesis plausible, one may not infer from the past correlation in the observed sample to a complete correlation in the total population."

If this is indeed an abductive inference, then "All A 's are B 's" should explain "All observed A 's are B 's." But, "All A 's are B 's" does not seem to explain why "This A is a B ," or why A and B are regularly associated (pointed out by Ennis, 1968). Furthermore, it is hard to see how a general fact could explain its instances, because it does not seem in any way to cause them.

The story becomes clearer if we are careful about what precisely is explained and what is doing the explaining. What the general statement in the conclusion explains are certain characteristics of the set of observations, not the facts observed. For example, suppose I choose a ball at random (arbitrarily) from a large hat containing colored balls. The ball I choose is red. Does the fact that all of the balls in the hat are red explain why this particular ball is red? No, but it does explain why, when I chose a ball at random, it turned out to be a red one (because they all are). 'All A's are B's' cannot explain why 'This A is a B' because it does not say anything at all about how its being an A is connected with its being a B. The information that "they all are" does not tell us anything about why this one is, except that it suggests that if we want to know why this one is, we would do well to figure out why they all are. Instead, all A's are B's helps to explain why, when a sample was taken, it turned out that all of the A's in the sample were B's. A generalization helps to explain some characteristics of the set of observations of the instances, but it does not explain the instances themselves. That the cloudless, daytime sky is blue helps explain why, when I look up, I see the sky to be blue (but it doesn't explain why the sky is blue). Seen this way, an inductive generalization does indeed have the form of an inference whose conclusion explains its premises.

In particular, 'A 's are mostly B's' together with 'This sample of A 's was obtained without regard to whether or not they were B's' explains why the A's that were sampled were mostly B 's.

> Why were 61% of the chosen balls yellow?
>
> Because the balls were chosen more or less randomly from a population that was two thirds yellow, the difference from 2/3 in the sample being due to chance.

Alternative explanation for the same observation:

> Because the balls were chosen by a selector with a bias for large balls from a population that was only one third yellow but where yellow balls tend to be larger than non yellow ones.

The frequencies in the larger population, together with the frequency-relevant characteristics of the method for drawing a sample, explain the frequencies in the observed sample.

What is explained? In this example, just the frequency of characteristics in the sample is explained, not why these particular balls are yellow or why the experiment was conducted on Tuesday. In general, the explanation explains why the sample frequency was the way it was, rather than having some markedly different value. If there is a deviation in the sample from what you would expect, given the population and the sampling method, then you have to throw some Chance into the explanation (which is more or less plausible depending on how much Chance you have to suppose).

How are frequencies explained? Observed frequencies are explained by giving a causal story that explains how the frequencies came to be the way they were. This causal story typically mentions both the population frequency and the method of drawing the sample.

Unbiased sampling processes tend to produce representative outcomes; biased sampling processes tend to produce unrepresentative outcomes. This "tending to produce" is causal and supports explanation and prediction. A peculiarity is that characterizing a sample as "representative" is characterizing the effect (sample frequency) by reference to part of its cause (population frequency). Straight inductive generalization (carrying the sample frequencies unchanged to the generalization) is equivalent to concluding that a sample is representative, which is a conclusion about its cause. Straight inductive generalization depends partly on evidence or presumption that the sampling process is (close enough to) unbiased. The unbiased sampling process is part of the explanation of the sample frequency, and any independent evidence for or against unbiased sampling bears on its plausibility as part of the explanation.

If we do not think of inductive generalization as abduction, we are at a loss to explain why such an inference is made stronger or more warranted, if in collecting data we make a systematic search for counter-instances and cannot find any, than it would be if we just take the observations passively. Why is the generalization made stronger by making an effort to examine a wide variety of types of A's? It is made stronger because the failure of the active search for counter-instances tends to rule out various hypotheses about ways in which the sample might be biased, that is, it strengthens the abductive conclusion by ruling out alternative explanations for the observed frequency. If we think that a sampling method is fair and unbiased, then straight generalization gives the best explanation of the sample frequencies. But if the sample size is small, alternative explanations, where the frequencies differ, may still be plausible. These alternative explanations become less and less plausible as the sample size grows, because the sample being unrepresentative due to chance becomes more and more improbable. Thus viewing inductive generalizations as abductions shows us why sample size is important. Again we see that analyzing inductive generalization as abduction show is how to evaluate the strengths of these inferences.

I have argued that it is possible to treat every smart (i.e., reasonable, valid, strong) inductive generalization as an instance of abduction, and that analyzing inductive generalizations as abductions shows us how to evaluate the strengths of these inferences. I conclude that inductive generalizations derive their epistemic warrants from their natures as abductions.

If this conclusion is correct, it follows that mechanisms for inductive generalization must be abductively well-constructed, or abductively well-controlled, if they are to be smart and effective.

References

Aristotle- explanatory syllogism - prior analytics I think

Aristotle - 4 causes, *Physics* , bk. 2, chap. 3.

Bhaskar, R. (1981). Explanation. In W. F. Bynum, E. J. Browne, & R. Porter (Eds.), *Dictionary of the History of Science*. (pp. 140-142). Princeton University Press.

Copi, I. M., & Cohen, C. (1998). *Introduction to Logic* (Tenth Edition ed.). Prentice Hall.

Ennis, R. (1968). Enumerative Induction and Best Explanation. The Journal of Philosophy, LXV(18), 523-529.

Fann, K. T. (1970). *Peirce's Theory of Abduction*. The Hague: Martinus Nijhoff.

Flach - probably this volume

Harman, G. (1965). The Inference to the Best Explanation. Philosophical Review, 74, 88-95.

Hempel, C. G. (1965). *Aspects of Scientific Explanation*. New York: Free Press.

Josephson, J. R., & Josephson, S. G. (Eds.). (1994). Abductive Inference: Computation, Philosophy, Technology. New York: Cambridge University Press.

Kruijff - probably this volume

Lipton, P. (1991). *Inference to the Best Explanation*. London: Routledge.

Lucas, On the Modeling of Interactions in Diagnosis, *Artificial Intelligence* , forthcoming

Öztürk - probably this volume

Pearl - ??

Poole - probably this volume

Psillos - probably this volume

Salmon, W. C. (1990). *Four Decades of Scientific Explanation*. Minneapolis: University of Minnesota Press.

Wallace, W. A. (1972). *Causality and Scientific Explanation, Volume 1, Medieval and Early Classical Science*. Ann Arbor: University of Michigan Press.

Wallace, W. A. (1974). *Causality and Scientific Explanation, Volume 2, Classical and Contemporary Science*. Ann Arbor: University of Michigan Press.