ELSEVIER

New Ideas in Psychology 25 (2007) 16–36

NEW IDEAS IN PSYCHOLOGY

www.elsevier.com/locate/newideapsych

Rationality and reasoning with metaphors

Bipin Indurkhya

International Institute of Information Technology, Gachibowli, Hyderabad 500 032, India

Available online 11 December 2006

Abstract

Traditional accounts of rationality typically preclude metaphorical reasoning. We review research that has highlighted the pervasiveness of metaphors in creative problem solving, jurisprudence and history of science, and argue that any account of rationality must explicitly acknowledge the ontology of representation and include an ontology-changing mechanism. From these considerations, we present an interaction-based view of cognition and examine the problem of rationality in its terms. We argue that rationality is closely related to the attitude of a cognitive agent towards incoherency—by which we do not mean internal inconsistency but operational incongruity in the external world. We conclude that though rationality permits a change of ontology, it requires a healthy respect towards the autonomous structure of the environment.

Keywords: Cognition; Creativity; Interaction; Metaphor; Ontology; Rationality; Reasoning

1. Introduction

Metaphors and analogies are often used as tools of reasoning. Sometimes they are so subtle that we are not even aware that a metaphor or analogy is at work in providing the justification for reaching a certain conclusion. At other times, they are quite explicitly evoked and stand out larger than life. Examples of both can be found in everyday reasoning as well as in academic discourse in science, philosophy, religion, etc. Lakoff and Johnson (1980) have done an excellent job at demonstrating how pervasive subtle metaphors are in everyday reasoning. For instance, the conclusion that you can "save" time by taking the subway to work instead of walking derives from the "time is money" metaphor that most of us take for granted. On the other hand, when a government

E-mail address: bipin@iiit.ac.in.

0732-118X/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.newideapsych.2006.10.006

announces that it is declaring a "war" on the nation's budget deficit and asks the citizens to offer sacrifices and put up with discomforts, the underlying metaphor is blatantly obvious—the "sacrifice" is increased taxes and the "discomfort" is reduced services. (See also Kempton, 1987, for a study of how people have different underlying metaphors for the way home heating works, and how these metaphors influence the way these people adjust their thermostats.) In a beautifully written essay titled *The Myth of Metaphor*, Turbayne (1962) points out that Descartes' "world is a machine" metaphor has worked its way so subtly into all of the Western sciences since his time that most scientists and non-scientists alike are no longer consciously aware of it.

We should emphasize that the issue here is the role of metaphor in constructing or laying out an argument, and not whether the argument is valid or not. For instance, Francesco Sizzi is reported to have used the following argument against Galileo's discovery of the satellites of Jupiter:

> There are seven windows in the head, two nostrils, two eyes, two ears, and a mouth; so in the heavens there are two favorable stars, two unpropitious, two luminaries, and Mercury alone undecided and indifferent. From which and many other similar phenomena of nature, such as the seven metals, etc., which it were tedious to enumerate, we gather that the number of planets is necessarily seven. (Quoted in Newman, 1956, p. 822.)

This argument clearly employs some very overt metaphors, though its conclusion was proven to be otherwise later. It seems that metaphors (in the rest of this essay my term 'metaphors' will subsume analogies as well) are quite commonplace in discourse, and sometimes they are there not merely for convenience but are the lifeblood of the argument because only they provide the necessary justification for accepting (or not accepting) a certain conclusion. We should emphasize here that this characterization of metaphor is rather broad as it subsumes the so-called conventional metaphor championed by Lakoff and his colleagues (Lakoff, 1987; Lakoff & Johnson, 1980) and the unconventional or novel metaphors that we have focused on elsewhere in our earlier work (Indurkhya, 1992a, 1992b, 2006).

Given how ubiquitous metaphors are, one might wonder: when is it rational to accept a conclusion derived from a metaphor? We may consider an extreme position: where the only justification that can be offered in support of a certain assertion comes from a metaphor, the assertion must necessarily be judged irrational—or, in other words, metaphors must be banned from all rational discourse. Though such a conclusion may seem rather radical, we should emphasize that the traditional accounts of rationality tend to not consider metaphorical reasoning at all (Harman, 1995). However, this extreme view precludes many instances of everyday as well as scientific discourse from being characterized as rational. In fact, some historians and philosophers of science would argue that in this traditional view all science, except perhaps mathematics, would have to be deemed irrational. For metaphors have been found to play a crucial and indispensable role in the growth and development of various sciences (Boyd, 1979; Gerhart & Russell, 1984; Hesse, 1966, 1974, 1980; Jones, 1982; Kuhn, 1979). It has even been argued that the structures of mathematics emerge from an interactive cognitive process that is similar, if not the same, to the one underlying metaphor (Lakatos, 1976; Lakoff & Núñez, 2000; Mac Lane, 1986). Thus, it seems that the proper answer to the question must lie somewhere in between: sometimes metaphors can be used to derive rational conclusions, sometimes not.

But then we must ask: What makes metaphors rational at some times but not at others? Addressing this question, as we will be doing in the rest of this essay, turns out to have a reciprocal effect. On one hand, considerations of rationality qualify when reasoning from metaphors is acceptable. On the other hand, certain metaphorical conclusions that seem quite reasonable—sometimes even necessary to maintain a rational worldview—have strong implications for rationality itself. We will be examining both of these issues here.

This essay is organized as follows. In the next section, I distinguish between two modes of metaphorical reasoning, which I refer to as monotonic and nonmonotonic reasoning, and point out their characteristics. In Section 3, I discuss monotonic reasoning with metaphors at some length, and then point out how considerations from rationality qualify it. In Section 4, I take up nonmonotonic reasoning with metaphors. I present some examples to show how it forms an indispensable part of rational discourse, and then proceed to point out the implication it has for rationality. In Section 5, I present an interaction-based view of cognition, and discuss how rationality might be viewed within it in a way that incorporates nonmonotonic reasoning with metaphors.

2. Two modes of reasoning with metaphors

Though I presented a few examples of metaphorical reasoning in the introduction, it would be useful to explicitly clarify what metaphorical reasoning means and to introduce a couple of terms before proceeding further. In metaphorical reasoning, two distinct domains are involved, and the implicative complex (or the knowledge) of one domain is used to reason (or to make an assertion) about the other domain. The first domain is referred to as the *source* and the second domain, the one being reasoned about, as the *target.* In the examples I presented in the introduction, time, reducing the budget deficit, the world, and the Solar system are the targets of their respective metaphors, and money, waging war, a machine, and the human head are their respective sources.

Note that I did not say anything about *similarities* between the source and the target: any similarities between the two may or may not play a role in the reasoning process, as we will see shortly.

An instance of metaphorical reasoning usually falls in one of two classes. One class consists of what I have called *similarity-based metaphors* (Indurkhya, 1992a). Here, some similarities are noted between the source and the target, and from these known similarities, further similarities are stipulated. It is these stipulated similarities that allow the implicative complex of the source to generate assertions about the target. I refer to this kind of reasoning as *monotonic* because it extends the knowledge of the target monotonically: no conclusion that could be derived from the knowledge of the target alone is ever invalidated by the metaphor, but new conclusions are added.

An example discussed by Gentner and Jeziorski (1989) illustrates how monotonic reasoning with metaphor works. They point out how Carnot used an analogy from fluid flow to heat flow to hypothesize that the rate of heat flow might be proportional to the temperature difference just as the rate of fluid flow is proportional to the difference in levels. This conclusion was based on other known similarities between heat flow and fluid flow; for instance, that fluids flow from higher levels to lower levels and that heat "flows" from a higher temperature to a lower temperature.

In all instances of monotonic reasoning with metaphors, known similarities between the source and the target play a crucial role in the reasoning. They are seen to "justify" or provide a rationale for the assertion derived from the metaphor.

The other kind of metaphor, which I have referred to as *similarity-creating metaphor* (Indurkhya, 1992a), is *not* based on known similarities between the source and the target. In fact, in most interesting examples of such metaphors, there are no known similarities between the source and the target. However, after the metaphor is applied, if it is successful, it makes the target look similar to the source. In other words, the similarities are *created* by the metaphor. I refer to the reasoning with similarity-creating metaphors as *nonmonotonic* because it typically works by causing a radical shift in the way the target is conceptualized, and the inferences that are generated by the metaphor may be incompatible with the implicative complex of the target that existed before the metaphor.

I will forgo presenting examples of similarity-creating metaphors until Section 4, but let me just stress here that the known similarities between the source and the target—that is, the known similarities *before* the metaphor—do not play any role in nonmonotonic reasoning. This is the single crucial point that distinguishes it from monotonic reasoning with metaphors.

3. Rationality and monotonic reasoning with metaphors

In monotonic reasoning with metaphors—more commonly known as reasoning with similarity or reasoning by analogy—the known similarities between the source and the target play a crucial role because they are seen to provide a rational basis for stipulating further similarities. Of course, why there is such a justification is an issue that philosophers have grappled with for a long time without success. It is interesting to point out here that this problem is closely related to the Humean problem of the justification of induction (Hume, 1987). Indeed, some of the theories that seek to provide justifications for reasoning with similarities are corollaries of more comprehensive theories of induction (Carnap, 1962; Harrod, 1956; von Wright, 1965), and some of the arguments I am going to make here parallel the ones that I have made elsewhere concerning the rationality of induction (Indurkhya, 1990).

The attempts to provide a rational foundation to reasoning with similarities can be divided into two classes: logical and empirical. To provide a logical foundation, the goal is often to show that an inference derived from a similar source is more probable than, say, a random inference. (See also the discussion in Hesse, 1966, pp. 101–129) Many elaborate mathematical systems have been proposed that try to demonstrate this in one of two ways. In the inductive approach, the idea is to show that known similarities between the source and the target form confirming instances of the hypothesis that the source and the target are similar (Carnap, 1962, pp. 569–570; Harrod, 1956, pp. 123–127. See also Achinstein, 1963; Carnap, 1963; Hesse, 1964 for refinements of these theories). In the generalization approach, the hypothesis that known similarities imply inferred similarities is analyzed in terms of the scopes of the antecedent and consequent, and probabilities are assigned to the generalization in direct proportion to the scope of antecedent and in inverse proportion to the scope of consequent (Keynes, 1921, Chapter XIX, pp. 222–232; von Wright, 1965, pp. 134–136). This has the effect that a generalization inferring fewer similarities from more known similarities comes out as more probable than the one inferring more similarities from fewer known similarities.

Both these approaches are problematic, though (see Indurkhya, 1992a, pp. 323–327 for more details). The induction approach can be used to assign high probabilities to two contradictory conclusions made from two different sources that share the same set of similarities with the target. The generalization approach fails because it makes the existence of the source superfluous—the generalization probabilities are assigned entirely based on the scopes of the antecedents and consequents.

To provide an empirical foundation for reasoning with similarities, the usual approach is to demonstrate the pragmatic success of the reasoning process both in psychological experiments and in real-world problem-solving situations (Clement & Gentner, 1991; Gentner, 1989; Gick & Holyoak, 1980, 1983). But each of these approaches is also fraught with controversy (Indurkhya, 1992b). The psychological experiments all satisfy the precondition that the "right" source—that is, a source that can successfully lead to the solution of the problem—is known to the experimenter. In other words, what we can conclude from these experiments is that reasoning from a source that is similar to the target can be effective provided that the "right" source is given. For real-world problems, when no such oracle is available for consulting, it is far from clear whether reasoning with similarities brings about any pragmatic success.

Most efforts to analyze real-world problem-solving situations take full benefit of hindsight, and often end up committing the fallacy of mixing up the before-metaphor situation with the after-metaphor situation (Schön, 1963). When a particular metaphor is successful in solving a real-world problem, it always ends up making the post-metaphor target seem similar to the source—we will see some examples of this in the next section—but that does not necessarily imply that the similarities were there before the metaphor, or that the similarities were the reason for choosing the source in the first place.

To appreciate the futility of seeking a new or distinctive rational basis for reasoning from similarities alone, one must look at some examples where stipulating further similarities from known similarities results in an unwarranted conclusion. We come across such examples everyday in the media, in political rhetoric, and in advertisements. We are not always aware of them. Sometimes, we already believe in the conclusion that is being derived from the similarities (so that no justification is really needed, as far as we are concerned). Sometimes, the analogy is so subtle that we find them psychologically compelling. It is only after we subject each argument from similarities to a careful logical analysis (Nosich, 1982), we become aware of the abuses of analogy.

Elsewhere (Indurkhya, 1992b) I have presented some examples of what I called the "dark side" of reasoning from similarities, because proponents of analogical reasoning conveniently choose to ignore them. For instance, Robbins (1991) pointed out similarities between the 1991 debate in the US Congress to authorize the President to use force against Iraq, and the 1939 debate when the US was considering whether to intervene in Europe. This article appeared when the multinational forces led by the US stood poised to attack Iraq in a bid to liberate Kuwait, and there was a world-wide debate on whether military force should be used. The author claimed, "The similarity is more than superficial," but no single explanation was offered as to how. Perhaps the author felt that the similarities between the arguments used in the Congress alone warranted the conclusion that the 1991 situation in the Middle East is similar to the 1939 situation in Europe in other respects as well.

This is where Weitzenfeld's (1984) approach to reasoning from similarities is most enlightening. He pointed out that there is always a second-order generalization, namely that the similarities in the known respects *determine* similarities in some unknown respects, that underlies any argument from analogy. To use the terminology of Stephen Toulmin (1964), it is the second-order generalization that acts as the *warrant* of the argument from similarities.

There are two major consequences of Weitzenfeld's approach. One is that an argument based on similarities, and similarities alone, is not justified at all in any sense. The other is that the justification for an argument from analogy depends directly on the degree of rationality one assigns to the implicit second-order generalization. For example, to hold that similarity in nationality of two persons determines similarity in their native language seems quite reasonable, given the background general knowledge about languages people speak in different nations. But to maintain that similarities in make, model, and year of two automobiles determine similarity in their color seems groundless (unless there is some background knowledge connecting the make, model and year of an automobile with the color).

Thus, the problem of finding rationality in reasoning from similarity turns out to be no different from the problem of finding rationality of any hypothesis. Similarities alone do not add the least bit of justification to an argument that cannot stand on its own—they act only as a cognitive smokescreen to hide the background assumption (the second-order generalization) so that the proponent of the argument does not have to make an honest effort at justifying or defending it. This background assumption is the only warrant for an argument based on similarities, and ought to be explicitly justified, if it is not obvious already. Failing to do so turns similarity-based reasoning into a cognitive blind spot that can delude us into accepting certain conclusions as rational, when, in fact, they are not justified in any sense at all.

4. Nonmonotonic reasoning with metaphors

Nonmonotonic reasoning with metaphors is characterized by metaphors that are not based on existing similarities between the source and the target; rather, if successful, they *create* the similarities between the two. I am going to introduce this kind of reasoning with some examples drawn from creative problem solving, jurisprudence and the history of science.

4.1. Creative problem solving

Perhaps the most classic example of nonmonotonic reasoning with metaphors is a case discussed by Schön (1963, pp. 74–76). A product-development team was given the task of improving the performance of synthetic fiber paintbrushes, which left the painted surface gloppy and uneven, as opposed to the smooth and continuous finish of natural fiber paintbrushes. Initially, their understanding of the painting process, and the role of paintbrush in it, was based on seeing it as a smearing process: the paint sticks to the fibers when the brush is dipped in it, and when the brush is moved back and forth on the surface, some of this paint sticks to the surface. From this perspective, when the product development team compared the performances of the natural fiber brush and the synthetic fiber brush, they could find no differences. Yet the appearance of the painted surface was different.

The breakthrough came when a member of the team had a flash of insight suggesting that the painting process be viewed as pumping. The concept of painting, as well as the role of a paintbrush in it, was completely transformed. Now the paint is sucked up in the space between the fibers by capillary action as the brush is dipped in the can of paint, and, when the brush is moved back and forth on the surface, the increased pressure pumps out the paint through the space between the fibers, generating a thin spray at the place where the fibers bend away from the surface. When the performances of the natural fiber paintbrush and the synthetic fiber paintbrush were compared from this perspective, it was found that the natural fibers formed a smooth curve when the paintbrush was pressed against a hard surface like a wall, while the synthetic fibers bent at a sharp angle. This observation led to a number of innovative suggestions for fabricating gradual-bending synthetic fibers, some of which resulted in bristles that curved smoothly; the paintbrush made from these fibers did, indeed, produce a smooth painted surface.

There are a few important points that must be noted in this example. First of all, the two different ways of viewing the painting process structure it in radically different ways. The very causal structure that determines how the paint sticks to the brush after it is dipped in the can, and how it is coated on the surface when the brush is moved back and forth across it, is different in each case. Moreover, the primitives underlying the structures are different. For example, the space between the fibers plays no role in the structure of *painting-as-smearing* model*,* but is a key factor in *painting-as-pumping* model.

There are two other crucial points. The flash of insight about painting as pumping was not triggered by some existing similarities between the painting process and the pumping process. This is why the identification was shocking at first—a characteristic that accompanies all creative but insightful metaphors. Secondly, the similarities between painting and pumping were *a result* of the metaphor. That is, after a different ontology for painting was created, *then* there were similarities between painting and pumping. Schön (1963, 1979) was most vehement about both these points.

There are several other such examples discussed in the literature on creative problem solving (Gordon, 1961; Koestler, 1964, pp. 199–207). Gordon (1961), for instance, proposed a problem-solving mechanism called ''making the familiar strange,'' to be used when the familiar perspective on the problem does not help in reaching a solution. The idea here is to view the familiar problem (after having unsuccessfully tried to solve it from the conventional perspective and by familiar techniques) ''strangely'' by juxtaposing it with strange domains. The juxtapositions initially seem quite crazy, but then one of them might lead to an illuminating insight that can prove to be the key to the solution. This mechanism is, in fact, quite the opposite of reasoning from similarities, which looks at the most similar source to import hypotheses into the problem domain. According to Gordon's view—a view that is supported by many case studies provided there—any attempt to reason from a similar source would only lead to mundane and not particularly insightful hypotheses. This is because a similar domain is not going to change the perspective on the problem domain.

Thus, we see that nonmonotonic reasoning with metaphors is an indispensable aid to creative problem solving. Almost always, we develop ontologies and perspectives on objects and situations in certain contexts. With repeated encounters, we become familiar with these perspectives, and they work well in similar contexts. However, when we encounter the same objects and situations in a radically different context, and we want to achieve certain ends, this becomes a problem and the familiar perspectives may no longer be adequate for solving it. Moreover, the context may be so different that no monotonic extension of the familiar perspective would do. What we need is a completely different perspective—as if we were encountering the object or the situation for the very first time. Techniques like "making the familiar strange" are really aimed at recovering our cognitive innocence, at least in part. But then any view of rationality ought to include them as well, for, otherwise, we would have no way to allow a shift from painting-as-smearing to painting-as-pumping in a rational discourse.

4.2. In jurisprudence

Legal reasoning is fraught with reasoning from similar past cases. In common law systems, as in England, previous cases are often the only sources of law. Even in statutory law systems, as in most of the continental Europe, past cases serve as important beacons for guiding how statutory predicates should be applied to the particulars of any given case. For a large part, legal reasoning proceeds along familiar lines—familiar to lawyers, that is—where the particulars of a given case suggest certain legal concepts and categories that might be relevant to it, and the past cases provide a constellation of concepts and categories from which different arguments can be made as to how the given case ought to be decided. (I must emphasize that I am by no means implying that such arguments are always clear-cut, or that the past cases often suggest an unambiguous resolution of the given case.) However, there are situations when the judges and the lawyers create new concepts and categories in order to apply a precedent to the facts of a given case, or to distinguish between them.

One such case concerns a concert violinist, Ernest Drucker, who was employed by the Metropolitan Opera (*Drucker* vs. *C.I.R.,* 715 F.2d 67, 1983). Drucker maintained a studio at home where he practiced regularly, and claimed a home-office deduction on his income tax returns. Internal Revenue Service (IRS) disallowed the home-office deduction, noting that under the tax statute, home-office deductions can be claimed only if the home office is the "principal place of business" for the taxpayer, and arguing that Drucker's principal place of business was the concert hall and not his home studio.

Drucker petitioned the Tax Court, where the majority opinion sided with the IRS. They argued that Drucker was an employee of the Metropolitan Opera, and consequently, was in the business of being a concert musician for the Met. Therefore, his principal place of business was the "focal point" of his activities as a musician for the Met, namely the concert hall. I must point out here that the Tax Courts had used the "focal point" test consistently in the past to determine the taxpayer's principal place of business. According to this test the principal place of business is the place where the most visible part of the taxpayer's business activity takes place; for example, the principal place of business for a teacher is the classroom, and not her office where she may prepare lecture notes and grade papers—even though she may spend a significantly more time there than in the classroom.

On appeal, the appellate courts reversed this decision. However, the majority opinion did not seem to want to disturb the focal-point test, nor did they seem to wish to deny that Drucker was in the business of being an employee of the Met. So they made a distinction between the principal place of business of an employer and the principal place of business of an employee, and argued that the two do not always have to coincide. The majority opinion held that the focal point of *Drucker's* business was his home practice, and, consequently, his principal place of business was his home studio, though the principal place of business of his employer, the Met, was the concert hall.

The important thing to notice here is that this distinction between the principal place of business of an employee and the principal place of business of an employer was made here for the first time, it seems, solely to justify the decision for Drucker without overruling the precedents. This distinction created a new ontology, on which basis the facts of Drucker's situation and tax responsibilities could be distinguished from the other precedents.

In another example, the Tax Court allowed a home-office tax deduction for an anesthesiologist, Nader Soliman, who maintained an office at home where he did administrative work like billing his patients (*Soliman* vs. *C.I.R.,* 94 T.C. 20). (This decision was later affirmed by appellate courts, 935 F.2d 52, 4th Cir. 1991, but reversed by the Supreme Court, 113 S.Ct. 701.) However, there was a preceding case where an emergency care physician, Stanley Pomarantz, was denied a home-office deduction in similar circumstances (*Pomarantz* vs. *Commissioner,* 867 F.2d 495, 9th Cir. 1988). In order to distinguish the Soliman case from the Pomarantz case, the majority opinion on the Tax Court described the similarities between the two cases as "superficial" and argued that the facts of the two cases are quite different because Pomarantz spent an "insubstantial" amount of time at his home office (8–15%), whereas Soliman spent a "substantial" amount of time at his home office (slightly over 30%). Here a new ontology was created by introducing the category "substantial" under which the facts of the two cases were distinguished.

Thus, we see that legal reasoning sometimes requires the creation of new categories, or a change of ontology, especially when unconventional and creative arguments are being made. The created categories can make the facts of two different cases seem similar, or they can distinguish between cases that would otherwise seem similar.

4.3. The history of science

Historians of science have pointed out that in many instances the development of science was crucially affected by certain analogies or models (Canguilhem, 1963; Gruber, 1978; Hesse, 1966; Koestler, 1964; Libby, 1922; Miller, 1978; Rothbart, 1984, pp. 611–612). While some of these analogies were guided by then-known similarities between the source and the target, as in the example of Carnot's use of the fluid-flow analogy to reason about the heat-flow that was mentioned earlier, in many cases the source was applied to the target even though there were no known similarities between the two phenomena at that time. These latter kinds of analogies, which have also been referred to as "pre-theoretic models" or "material models," usually work as follows. There is a target phenomenon that is poorly understood, and a source that is more concrete and well understood. But there is no good connection or transference between the two phenomena that would form the basis for the source to be meaningfully applied to the target. There is only some vague idea in the mind of the scientist. Sometimes, there is not even a vague idea but an emotional drive that keeps the scientist seeking some connection. However, once the connection is found, if it is found, the model starts to become entrenched. If we compare the states of the source and the target before and after the model became entrenched, the creation of similarities can be seen quite starkly. The source is seen to create an ontology for the target phenomenon, where none existed before.

I present two examples here to illustrate these points. Gruber (1978), on analyzing Darwin's notebooks that he kept while working on his celebrated theory of evolution, noted that the image of an irregularly branching tree kept on recurring in his thoughts, and may have served as the primal metaphor for his natural selection principle. Gruber's account cites several excerpts from Darwin's notebooks to show that it was not some existing similarities between the tree image and whatever was known about the evolution at that time that kept Darwin searching for what other similarities might be found. On the contrary, it was an emotional drive that kept up his intellectual commitment to articulating an account of evolution that matched the tree image. Darwin was not searching for similarities, he was formulating them. This process was actually fraught with several problems that appeared to invalidate the model, but Darwin's intellectual commitment to the idea kept him from throwing away the image altogether. Instead, he improvised and sought different ways to render the image meaningful.

In another example, Solomon (1992) points to Wegener's (1966) use of the iceberg analogy to suggest the theory of continental drift. Solomon notes that the theory of continental drift was not widely accepted by the geologists until the mid-1960s—after the discovery of seafloor spreading, which led to the development of the modern theory of plate tectonics—and was considered quite revolutionary when Wegener first proposed it in 1915. Here again, the analogy was not based on the known similarities between the continents and the polar icebergs. In fact, Solomon explicitly notes: "[T]he situations of icebergs and continents [were] not shown to be sufficiently similar to support the conclusions drawn" (1992, p. 447). I should also point out that Solomon attributes Wegener's strong attraction to the drift theory partly to the fact that he was not trained as a geologist, and so he was not indoctrinated with the contractionist theory—the leading geophysical theory of continents at that time. The reader should be able to see a connection here with Gordon's *making the familiar strange* and Schön's story of the paintbrush-as-a-pump metaphor. In order to conceptualize a phenomenon differently, and to see it from a different ontology, it is necessary that one be not attached too strongly to the conventional or familiar way of looking at it.

Thus, we see that nonmonotonic reasoning with metaphors—reasoning that is not triggered by some existing similarities between the source and the target, but is constitutive of them—plays a significant role in the development of science, and must not be excluded from rational discourse.

4.4. Implications for rationality

The preceding examples suggest that metaphors can radically alter the ontology of the target, thereby creating a new perspective on it that includes information not contained in the old representation. Moreover, this new information often contains the key to resolving a problematic situation that seems impossible from the familiar perspective. This role of metaphors, though it is essentially non-algorithmic and heuristic-driven, should not be excluded from an account of rationality, because, otherwise, one would have to admit irrationality as a key ingredient of major areas of human thought and discourse.

The whole problem boils down to incorporating an ontology-changing move into rationality. This move should allow introducing new concepts and categories to describe an object or a situation. Moreover, the new concepts and categories are not always added to the old concepts and categories, but may replace them—in the context of the current problem, if not always. Thus, any account of rationality must

1. explicitly take into account the ontologies of representations,
2. contain mechanisms for changing ontologies, and

3. allow different representations of the same object or phenomenon to coexist, and be able to weigh their relative merits in any given context.

Allowing an ontology-changing move into rational discourse is not as simple as it may appear at first. Though a rational cognitive agent may initiate a change of ontology, the new ontology may not always lead to the desired effect. The following examples illustrate this problem:

1. Schön noted that the group working on the synthetic-fiber paintbrush problem also considered "painting as masking a surface" metaphor, with its corresponding change of ontology, but it led to no useful insight (Schön, 1979, p. 259).
2. Yolanda Baie, a food stand operator and owner, petitioned to have the kitchen of her house, where she prepared food sold at the food stand, qualify as a "home office" for the purpose of tax deduction (*Baie* vs. *C.I.R.,* 74 T.C. 105). One of her arguments was that her kitchen was a "manufacturing facility" for her business. The judges, while finding the argument "ingenious and appealing," ruled it "insufficient" nonetheless.
3. John Casti (1989, pp. 7–10) comments on the fate of Immanuel Velikovsky's theory as outlined in *Worlds in Collision,* which hypothesized Earth's encounters with a large comet expelled from Jupiter and provided explanations for many biblical events. Velikovsky's theory proposes a revolutionary new ontology for understanding the Solar system. But it has not been accepted by the scientific community.

Thus, the problem is to somehow allow creative changes of ontology, but to rule out arbitrary changes of ontology; or, at least, to constrain which inferences are rationally acceptable from the changed ontology. (Palmer, 1992, pp. 62–64, raises the same problem and proceeds to address it in the rest of his book.) In the next section, I outline a view of cognition that addresses this problem.

5. Cognition and rationality

The problem of rationality mentioned at the end of the preceding section is remarkably parallel to a problem of creative metaphors: namely, how can some metaphors create new and useful perspectives on the target, whereas many others simply do not work? It also parallels a problem of cognition: how can a cognitive agent view an object (or an event, or a phenomenon) in many different ways, but not in any arbitrary way? As I have discussed both these latter problems at length elsewhere (Indurkhya, 1992a) by proposing a view of cognition as interaction and articulating a theory of metaphor within it, I will attempt here to address the problem of rationality within the same framework.

5.1. An interaction-based view of cognition

It will be useful at first to lay down the rudiments of the interaction view of cognition, which has been influenced by the works of Ernst Cassirer (1955), Nelson Goodman (1978), and Jean Piaget (1971a, b, 1977). Cognition is seen as an interaction between the cognitive agent and the external objective world; the agent and the environment are deemed to play an equal role in determining cognitive structures, which include scientific theories.

The key features of our view are as follows. We maintain a sharp distinction between an object (or an experience, or a domain, or an environment) and its representation (or description) and argue that an agent can cognitively access an object only by interacting with it in some way. Further, we claim that in order to have purposeful interactions with the world, the cognitive agent must use some internal states or structures. We refer to these internal states as "representations."

I must emphasize here that our use of the term "representation" is somewhat different from the way it has generally been used in artificial intelligence and philosophy of mind. Our "representation" is a meta-level term that refers to the internal states of the cognitive agent. We do not wish to imply that the cognitive agent has to be consciously aware of the representations or what they mean (how they are connected to the world). We do not even aim to imply that the representations have to be symbolic, though, talking about the internal states from a meta-level, we can certainly assign symbols to them. (See Indurkhya, 1992a, 1992b, pp. 135–151, for an example where representations are internal states of a simple cognitive agent. Therefore, arguments *à la* Brooks, 1991, do not automatically run contrary to our position.)

We distinguish between the ontology and the structure of a representation. The ontology is the set of primitives in terms of which the object is described, and the structure is the particular way in which the primitives are interrelated in the description. When the cognitive agent has internal states, the set of states constitutes the ontology of the representation, and the state transition functions make up the structure. An important observation to make here is that every structure presupposes an ontology.

When a cognitive agent interacts with an object, the result is a representation of the object that is internal to the agent. This representation is determined in part by the agent and in part by the autonomous nature of the object. The agent gets to play a role in giving an ontology to the object, a process which may be partly constrained by the agent's own biological structure. The object gets to play a role by endowing this agent-determined ontology with a structure. Thus, the structure of the object always seems autonomous (external) to the agent, although the agent can choose a different ontology, and thereby affect the structure of the object as internally represented.

To elaborate on this representation-building process, we need to make an explicit distinction between its two poles. On one end are the concepts, which are the building blocks of representation, and on the other end are the sensory stimuli, which provide the raw data out of which representations are formed. Both these ends are crucial in our view of the representation-building process, in keeping with Nelson Goodman's rephrasing of the Kantian dictum: "Although conception without perception is merely *empty,* perception without conception is *blind* (totally inoperative)" (Goodman, 1978, p. 6, emphasis original). We refer to these two components of the representation-building process as concept networks and sensorimotor datasets, respectively.

Concept networks are structured networks of uninterpreted symbols and relations between them. The theory of pumping, the theory of iceberg drift, geometric concepts like "triangle" and theorems relating them are all examples of concept networks, as long as no interpretation is attached to them, meaning that the theories are not applied to the world. I must hasten to add that we do not wish to suggest that these theories are devoid of any meaning for the cognitive agent. Obviously, all these theories were developed to explain certain phenomena and are, therefore, representative of them. All we mean to say is that at the meta-level, in order to explicitly talk about the representation building process, we can sever the meaning connection and refer to the now uninterpreted theories as concept networks.

A concept network may be genetically inherited, or it can be learned from the past interactions with the environment. The ontology of a concept network refers to the set of concepts, and the structure refers to the statements and theorems. Thus, if we write a theory in a formal system, its language is the ontology and the set of sentences of the theory forms the structure. Notice that the semantic content of a concept network manifests itself in its structure: the way the concepts are interrelated to each other reflects the structure of the phenomenon to which the concept network usually applies.

Sensorimotor datasets are structured sets of stimuli received from the environment, such as the set of images formed on the retina as a cognitive agent watches the process of painting. Sensorimotor datasets provide the raw input to the representation-building process. It is useful to apply the distinction between ontology and structure to the sensorimotor datasets also. The set of raw sensory stimuli, such as the set of retinal cell excitations, together with the set of correlations between different stimuli that the cognitive agent is capable of noticing, constitute the ontology of the sensorimotor dataset. As our perceptual apparatus creates the sense data, it suggests that the ontology of the sensorimotor dataset is determined by the perceptual and motor organs of the cognitive agent. However, as noted by Rosch (1978, p. 28), reality, as it presents itself to us in the form of sensory stimuli, is already highly structured: certain sets of stimuli occur together, while certain others preclude one another. This specific pattern of correlations that actually exists between different stimuli makes up the structure of the sensorimotor dataset. Moreover, this structure of the sensorimotor dataset is determined by the mind-independent external world.

This distinction between the ontology (determined by the cognitive agent) and the structure of a sensorimotor dataset (determined by the external world) is crucial. An example would be helpful to clarify it further. Suppose I am standing in front of a tree. It is the biological structure of my eyes that determines what image is formed on my retinas, what contrast I see, what colors I see, etc. (Of course, it also depends on the object in front of my eyes, the lighting conditions, etc. But different objects, or the same object under different lighting conditions, can all be considered different stimuli in the sensorimotor dataset.) Similarly, as I walk towards the tree, it is the structure of my motor apparatus that determines what movements I can make. However, the effect of any such movement on changing one stimulus on my retina into another, possibly different, stimulus is determined by the structure of the external world. It is precisely in this sense that reality, while owing its form (as experienced by me) to my sensorimotor apparatus, nonetheless appears as an autonomous external entity to me. (This example may seem somewhat simplistic because our sensory organs already act as pre-processors of sensory information, and one could argue that the structure of this pre-processed sensory information is due, at least in part, to the structure of the sensory organ. However, as argued in Indurkhya, 1992a, pp. 179–187, in this situation, the account of the representation-building process being articulated here can be applied to the pre-processing of sensory information itself: in explaining the processing done even by simple sensory organs, one could posit a sensorimotor dataset with an agent-driven ontology and an environment-driven structure.)

Finally, it must be emphasized that a sensorimotor dataset includes only perceptual data; in particular, we do not allow information obtained by instruments to be a part of it. The reason is that when an instrument is being used, it ultimately provides us with some perceptual data, such as the deflection of a needle, an image seen in a photograph or through the eyepiece of a telescope, or squiggly lines on graph paper, and one needs to apply a theory (that is, carry out the representation-building process) to interpret the perceptual data as the current flowing through a wire, a planet up in the sky, or the activity of certain brain cells (see also Lakatos, 1978, pp. 12–19).

In building the representation of an object (or event or phenomenon), the cognitive agent applies a concept network to the object. Some of the time, perhaps most of the time, this process will be habitual—the cognitive agent may have biological or cultural preferences for applying certain concept networks to certain objects or phenomena in certain ways. Recall that all the examples of concept networks presented above have semantic contents, thereby suggesting that there are certain objects or phenomena in the world to which they conventionally apply. However, we can get a glimpse of the representation-building process when a concept network is applied to a new phenomenon (as in viewing painting as pumping), or when the phenomenon reveals unexpected behavior with respect to the structure of the concept network (as when the synthetic-fiber paintbrushes left the painted surface gloppy and uneven). Such situations reveal two principal mechanisms of representation building.

Projection, derived from the Piagetian notion of assimilation (Piaget 1962, 1977),[1] is used in applying a concept network to a phenomenon to which it does not normally apply. In this process, parts of the sensorimotor dataset of the object are grouped together, and labeled with the concepts from the concept network. In the grouping process, the sensorimotor dataset acquires an ontology as seen from the conceptual layer. However, the autonomous structure of the sensorimotor dataset asserts itself by endowing the conceptual ontology of the phenomenon with an autonomous structure too. In the context of the paintbrush-as-a-pump example, this means that though the cognitive agent can group the process of painting in various ways so as to see it as a pumping process (this process will also be constrained by the existing semantic content of the "pumping" concept network), whether this grouping satisfies the predictions made from the concept network (such as whether certain design of synthetic-fiber paintbrush will result in an even painted surface) depends entirely upon the autonomous nature of external reality.

Projection is non-algorithmic and heuristic-driven. The reason behind it is that an ontology is a prerequisite for an algorithm. For example, in seeing the paint-brush as a pump, there is no algorithmic way to arrive at the appropriate ontology. In fact, there is no algorithm to even decide whether an appropriate ontology exists. This is why in creative problem-solving situations, it is impossible to determine a priori which metaphors will work and which will not. Projection crucially involves what Harold Brown has called *judgment* (Brown, 1988, Chapter IV), and may reflect the cognitive agent's cultural background (both inherited and learned) as well as past individual experiences.

When the object, owing to its external autonomous structure, reveals some unexpected behavior—in other it words, it was not anticipated by the cognitive agent based on the structure of the concept network that was projected on the object—the cognitive agent

[1]The term "projection" is used, with a somewhat different meaning, for the first of the two phases of reflecting abstraction (see Piaget, 2001).

may choose to restructure the concept network so as to integrate the new behavior. This process is referred to as *accommodation* (Piaget, 1962, 1977).[2]

Unlike projection, accommodation can be quite algorithmic. This is because in accommodation, one stays within the same ontology, there is some specific structural statement (or set of statements) that is found to be incoherent with the environment, and what is done is to restructure the theory enough so that the incoherent structural statements are no longer included. (For some algorithms of accommodation, see de Kleer, 1986; Doyle, 1979.) Of course, accommodation is not always deterministic, for there may be more than one way to exclude the incoherent structural statements.

Although projection and accommodation act in consort in building representations, one can find examples where one or the other mechanism plays a predominant role. Such examples help to illustrate the representation-building process, and also the mechanisms of projection and accommodation themselves.

The system of latitudes and longitudes, and the division of day into hours, minutes and seconds are both examples of representations that have been built primarily by projection. The method Imre Lakatos (1976) has referred to as "monster-barring" in mathematical theorem proving can also be viewed as projection, because the mathematician is essentially redefining the concept (thereby changing the ontology) so as to exclude the "monster" examples. Note that projection is non-algorithmic and heuristic in all these cases. For instance, the discussion in Lakatos (1976) reveals how various new definitions of "polyhedron" were coined in order to exclude certain troublesome solids. These new definitions, however, were always based on noticing (or, perhaps, "creating") some new attribute of the solids so that the attribute would apply to the desired polyhedra but not to the monsters. Some very novel attributes were created in the process. Considering them, it seems impossible for an algorithm to generate all possible attributes of a solid that could be used as a basis for creating different ontologies.

Here is an important point about projection: even though the cognitive agent chooses a concept network (with an ontology and a structure) and groups the sensorimotor dataset, thereby giving it an ontology as seen from the conceptual layer, once a concept network is projected, the structure of the environment as seen through the projection is determined by the objective world. Once the lines of latitude and longitude are put in place, whether any two given places lie on the same latitude is no longer an arbitrary matter. Once a definition of polyhedron is accepted, it becomes an objective matter whether a certain solid is a polyhedron, and whether the numbers of its faces, edges, and vertices bear a certain relationship. Notice that once the ontology is fixed, algorithms can be devised to determine structure and check for structural coherency: once the system of latitudes and longitudes has been projected, we can come up with algorithms to decide whether two places have the same latitude or not.

On the other hand, geographical or political maps and the division of year into days are examples of representations formed primarily by accommodation. Here the ontology is kept fixed, and the structure of the concept network is adapted to reflect the structure of the environment as seen from the given ontology. Lakatos's (1976) method of revising a conjecture in the face of counterexamples is also an example of accommodation. Notice here the algorithmic nature of accommodation. Once the ontology of a geographical map

[2]Accommodation is related, though not identical, to reflection, the second phase of reflecting abstraction (Piaget, 2001).

is fixed, one can come up with an algorithm for taking measurements on the ground and transferring these measurements to the map (which corresponds to the concept network). In Lakatos's examples, existing attributes of geometric solids are used to limit the scope of the conjecture—a process that is very different from that of inventing new attributes as in case of projection.

I must emphasize that in accommodation, the cognitive agent also gets to play a key role in building the representation by fixing the ontology, so it is not completely an environment-driven process. Moreover, in map-making at least, there are instances where ontology itself becomes an issue, thereby making projection more dominant in the representation building process. The debate over what constitutes a "wetland" is a case in point (Mitsch & Gosselink, 1986, pp. 15–20).

Finally, I must note that the two-layered account of representation-building process presented here, in which projection and accommodation are at work between one concept network layer and one sensorimotor dataset layer, can be extended into a model where the process of integrating sense data into concepts proceeds through several layers of intermediate representations (Indurkhya 1992a, pp. 179–187).

6. Rationality in the interaction-based view

We can now address the problem of incorporating an ontology-changing move into rationality. Notice that in the foregoing account of cognition, the representation-building process pits the concept network, which is structured by the cognitive agent, against the sensorimotor dataset, which is structured by the external environment. As these structures are independent, one may not respect the other; we will refer to this characteristic as *incoherency* (or *coherency,* when the two structures do respect each other). We claim that considerations of rationality demand that incoherency, whenever it is found, be duly acknowledged and dealt with; however, they do not constrain the manner in which incoherency is dealt with; in particular, they do not favor either projection or accommodation as a primary mechanism for resolving incoherency. I will elaborate on this remark in the rest of this section.

Let us say that a certain theory has been developed to explain a phenomenon. In our terminology, the theory is a *representation* of the phenomenon. In order to explicitly talk about different theories applying to the phenomenon, and the theory applying to other phenomena, we sever the meaning connection between the representation and the phenomenon. Now the theory corresponds to our concept networks and the phenomenon becomes a sensorimotor dataset. Of course, the cognitive agent normally uses the theory to comprehend and interact with the phenomenon, so the structure of the concept network is really the structure of the phenomenon as seen from the conceptual layer.

Until some point in time, the theory made many correct predictions about the phenomenon. Then an event was found that contradicted the predictions made by the theory: the representation turned out to be incoherent. Rationality would demand that this contradiction be acknowledged, for it would certainly be quite irrational to simply ignore the available evidence and insist on the correctness of the theory. However, there are two ways in which the incoherency can be dealt with.

One alternative is to keep the same ontology, but use accommodation to alter the structure of the concept network: in other words, revise the statements of the theory so as to incorporate the troublesome observation (as well as all of the past observations that are relevant to the phenomenon). For instance, after noticing that synthetic-fiber paintbrushes do not produce as good-looking a painted surface as natural-fiber brushes do, one could try to account for it by noticing some difference between the natural and the synthetic fibers, while staying within the painting-as-smearing theoretical framework. Of course, any such explanation will be subject to further testing, and later observations may render it incoherent. In the synthetic-fiber paintbrush story, Schön observed that the members of the product development team noticed that the natural fibers all had split ends. Thinking that this explained the appearance of the painted surface, they tried to make synthetic fibers with split ends too, but it did not yield the expected result.

The other alternative is to use projection to alter the ontology of the phenomenon (Lakatos, 1976, 1978, pp. 16–17). This can be done by noticing new attributes and relations in the phenomenon, using them to re-categorize the objects and events, noticing correlations between these new categories, and creating a theory to express these correlations. Or, one can take a different theory—a theory that was not developed in the context of this phenomenon—and try to understand the phenomenon in terms of the new theory, as happened in the paintbrush-as-pump example. This process of projecting a new theory onto a phenomenon is constrained from both ends. On one end, the existing structure of the concept network, which reflects the structure of the phenomenon for which the theory was originally developed, resists arbitrary application; on the other end, the autonomous structure of the phenomenon, which is rooted in the external world, resists arbitrary categorization. For instance, in seeing painting as pumping, one has to respect the conceptual structure of "pumping" as well as the structure inherent in the actual process of dipping the paintbrush in a can of paint and moving the paintbrush back and forth across a surface. When this process of projection succeeds, if it succeeds at all, what results is a new ontology of the phenomenon with a structure that reflects the structure of the new theory.

The important thing to notice in projection is that each time the ontology is changed, the environment endows this ontology with a structure that is external, and hence appears objective, to the cognitive agent. And to stay within rational discourse, the cognitive agent must respect this objective structure. Though, of course, the cognitive agent may shift the ontology again, if not satisfied with the structure seen in terms of the previous ontology, but then the external world gets to structure this new ontology, and so the game goes on. Needless to say, the cognitive agent may resort to accommodation at any stage of the game by sticking to an ontology and altering the structure of the theory.

It is interesting to note here that projection gets to play a key role both when a revolutionary new theory with a radically different ontology is proposed—when new discoveries are made by projecting a concept network onto an "unsimilar" phenomenon—and also when hypothetical entities are proposed in order to save an old theory in face of contradictory evidence. (See also Brown, 1983, for an account of how Piagetian assimilation, on which our notion of projection is based, can be used as a rational discovery strategy.)

Though considerations of rationality alone do not dictate any hard and fast rules as to when accommodation should be favored over projection and vise versa, some general guidelines emerge if we examine carefully why a change in ontology sometimes becomes necessary. Cognition typically involves categorization: different sensory stimuli, made available to us by our perceptual apparatus, are grouped into concepts and categories. The world as seen from the conceptual level is considerably more simplified and structured than the one at the sensory level. This simplification is necessary to make us survive in an infinitely complex world with our finite and limited minds. However, an act of categorization invariably involves loss of information. (See also Hesse, 1974, pp. 48–51, for arguments leading to the conclusion that any classification necessarily results in a loss of verbalizable empirical information.) When a bunch of objects are placed in a category, their potential differences are overlooked; one might say that they are discarded by the process of cognition. Similarly, when two objects get put in different categories their potential common grounds are discarded. Though the simplified worldview that results from grouping the world one way makes it easier for the cognitive agent to interact with the world, the agent is thereby deprived of a horde of alternate worldviews.

The concept networks that a cognitive agent develops—over a lifetime or over generations—for various phenomena and objects in the environment reflect the agent's goals and predispositions in its past interactions with those phenomena and objects. The simplified worldviews created by these concept networks may make it very efficient for the cognitive agent to go about its everyday activities; sometimes they may even make a crucial difference in whether the cognitive agent can survive in the environment. However, as the environment changes (e.g., synthetic fibers become available), and the cognitive agent's goals and predispositions also vary (e.g., the need for synthetic-fiber paintbrushes arises), sometimes the information lost in the categorization leading up to the existing concept network becomes the key to realizing the agent's goals. In such a situation, projection becomes a crucial asset to cognition, for only it can bring about a change in the ontology and recover (part of) the lost information.

All this suggests that whenever the domain of activity requires some stability in ontology, as in science and jurisprudence, accommodation is to be generally preferred. An ontology-changing move (projection) should be used sparingly—only on those infrequent occasions when every effort has been made to address the problem within the existing ontology but without any success. For projection is cognitively expensive, requiring a mammoth effort to try to render the new theory meaningful in the context of the phenomenon, even before it can be ascertained whether the new ontology satisfies the desired goals. Yet if a type of problem is particularly hard to solve from the familiar point of view, a rational agent should realize that it may be because some crucial information is missing from the existing conception of it, and a new ontology may be called for. Forgetting this, the agent may fall into the rut of viewing the phenomenon always from the same old perspective, thereby shunning all creativity. This is why Gordon (1961) touts "making-the-familiar-strange" as a key mechanism of creative problem solving.

On the other hand, there are certain domains where the goal is to create new perspectives—to reveal alternate worlds that are lost in the familiar perspective—for no other reason but their own sake. Conceptual and modern art are a good example of such domains (Fox, 1982). In these situations, projection ends up playing a dominant role, for obviously, it is the only mechanism that allows new perspectives to be created.

When we consider changes in ontology, another issue to ponder is whether there are any criteria that identify certain ontology-changing moves rational and others irrational. There is obviously a difficulty in positing any such criterion a priori. This is because in projecting a new theory onto a phenomenon, one cannot determine beforehand what the new ontology of the phenomenon would be and what consequences it might have. It is for these reasons that the crux of creativity remains rather elusive.

It is tempting to connect rationality with the a posteriori success of the ontology-changing move. This would have the consequence that the three examples of failed ontology-changing moves presented earlier in Section 4.4 would be dubbed irrational. Taking this position, however, puts too much of a burden on the concept of rationality, making rationality synonymous with correctness or truth. If we carefully analyze each of the examples of Section 4.4 without considering whether the ontology-changing move was successful or not, it is hard to come up with any grounds to reject them: to view painting as masking a surface seems no more of an offbeat idea than to view it as pumping; Baie's claim that her kitchen was a "manufacturing facility" seems no more ad hoc than the court's argument that Drucker's principal place of business was not the same as his employer's; and Velikovsky's theory of the Solar system seems no more groundless than Wegener's theory of continental drift at the time each one was proposed.

Thus, rationality is not so much a characteristic of an argument that predicts success. Rather, it is an attitude of the cognitive agent that acknowledges incoherency, decides when and how ontology-changing moves are made, and displays a healthy respect for the autonomous structure of the world.

References

Achinstein, P. (1963). Variety and analogy in confirmation theory. *Philosophy of Science*, *30*(3), 207–227.

Boyd, R. (1979). Metaphor and theory change: What is metaphor a metaphor for? In A. Ortony (Ed.), *Metaphor and thought* (pp. 356–408). Cambridge: Cambridge University Press.

Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, *47*, 139–159.

Brown, H. I. (1983). Assimilation and discovery. *New Ideas in Psychology*, *1*(2), 89–97.

Brown, H. I. (1988). *Rationality*. London: Routledge.

Canguilhem, G. (1963). The role of analogies and models in biological discovery. In A. C. Crombie (Ed.), *Scientific change* (pp. 507–520). New York: Basic Books.

Carnap, R. (1962). *Logical foundations of probability* (2nd ed.). Chicago: University of Chicago Press.

Carnap, R. (1963). Variety, analogy, and periodicity in inductive logic. *Philosophy of Science*, *30*(3), 222–227.

Cassirer, E. (1955). In R. Manheim (Ed.), *The philosophy of symbolic forms: Vol. 1–3*. New Haven: Yale University Press (Original work published 1923–1929).

Casti, J. L. (1989). *Paradigms lost: Images of man in the mirror of science*. New York: William Morrow & Co.

Clement, C. A., & Gentner, D. (1991). Systematicity as a selectional constraint in analogical mapping. *Cognitive Science*, *15*, 89–132.

de Kleer, J. (1986). An assumption-based TMS. *Artificial Intelligence*, *28*, 127–162.

Doyle, J. (1979). A truth maintenance system. *Artificial Intelligence*, *12*, 231–272.

Fox, H. N. (1982). *Metaphors: New projects by contemporary sculptors*. Washington, DC: Smithsonian Institution Press.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou, & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199–241). London: Cambridge University Press.

Gentner, D., & Jeziorski, M. (1989). Historical shifts in the use of analogy in science. In B. Gholson Jr., W. R. Shadish, R. A. Neimeyer, & A. C. Houts (Eds.), *Psychology of science: Contributions to metascience* (pp. 296–325). London: Cambridge University Press.

Gerhart, M., & Russell, A. (1984). *Metaphoric process: The creation of scientific and religious understanding*. Fort Worth: Texas Christian University Press.

Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, *12*, 306–355.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, *15*, 1–38.

Goodman, N. (1978). *Ways of worldmaking*. Indianapolis: Hackett.

Gordon, W. J. J. (1961). *Synectics: The development of creative capacity*. New York: Harper and Row.

Gruber, H. E. (1978). Darwin's "tree of nature" and other images of wide scope. In J. Wechsler (Ed.), *On aesthetics in science* (pp. 121–140). Cambridge, MA: MIT Press.

Harman, G. (1995). Rationality. In E. E. Smith, & D. N. Osherson (Eds.), *Thinking: Invitation to cognitive science*, Vol. III (pp. 175–211). Cambridge, MA: MIT Press.
Harrod, R. (1956). *Foundations of inductive logic*. London: Macmillan.
Hesse, M. B. (1964). Analogy and confirmation theory. *Philosophy of Science*, *31*(4), 319–327.
Hesse, M. B. (1966). *Models and analogies in science*. Notre Dame, IN: University of Notre Dame Press.
Hesse, M. B. (1974). *The structure of scientific inferences*. Berkeley, CA: University of California Press.
Hesse, M. B. (1980). *Revolutions and reconstructions in the philosophy of science*. Bloomington, IN: Indiana University Press.
Hume, D. (1987). *A treatise on human nature*. Harmondsworth: Penguin (Original work published 1739).
Indurkhya, B. (1990). Some remarks on the rationality of induction. *Synthèse*, *85*, 95–114.
Indurkhya, B. (1992a). *Metaphor and cognition*. Dordrecht: Kluwer Academic.
Indurkhya, B. (1992b). Predictive analogy and cognition. In K. P. Jantke (Ed.), *Analogical and inductive inference* (pp. 214–231). Berlin: Springer (Lecture Notes in Artificial Intelligence 642).
Indurkhya, B., 2006. Emergent representations, interaction theory and the cognitive force of metaphor. New Ideas in Psychology, in press, doi:10.1016/j.newideapsych.2006.07.004.
Jones, R. S. (1982). *Physics as metaphor*. Minneapolis: University of Minnesota Press.
Kempton, W. (1987). Two theories of home heat control. In D. Holland, & N. Quinn (Eds.), *Cultural models in language and thought* (pp. 222–242). Cambridge: Cambridge University Press.
Keynes, J. M. (1921). *A treatise on probability*. London: Macmillan.
Koestler, A. (1964). *The act of creation*. London: Hutchinson.
Kuhn, T. S. (1979). Metaphor in science. In A. Ortony (Ed.), *Metaphor and thought* (pp. 409–419). Cambridge: Cambridge University Press.
Lakatos, I. (1976). In J. Worrall, & E. Zahar (Eds.), *Proofs and refutations: The logic of mathematical discovery*. Cambridge: Cambridge University Press.
Lakatos, I. (1978). In J. Worrall, & G. Currie (Eds.), *Philosophical papers, vol. 1: The methodology of scientific research programmes*. Cambridge: Cambridge University Press.
Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago Press.
Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
Lakoff, G., & Núñez, R. E. (2000). *Where mathematics come from*. New York: Basic Books.
Libby, W. (1922). Scientific imagination. *Scientific Monthly*, *15*, 263–270.
Mac Lane, S. (1986). *Mathematics: Form and function*. Berlin: Springer.
Miller, A. I. (1978). Visualization lost and regained: The genesis of the quantum theory in the period 1913–27. In J. Wechsler (Ed.), *On aesthetics in science* (pp. 73–102). Cambridge: MIT Press.
Mitsch, W. J., & Gosselink, J. G. (1986). *Wetlands*. New York: Van Nostrand Reinhold.
Newman, J. R. (1956). *The world of mathematics*, Vol. 2. New York: Simon & Schuster.
Nosich, G. N. (1982). *Reasons and arguments*. Belmont, CA: Wadsworth.
Palmer, S. (1992). *Human ontology and rationality*. Aldershot, England: Avebury.
Piaget, J. (1962). In C. Gattegno, & F. M. Hodgson (Eds.), *Play, dreams and imitation in childhood*. New York: W.W. Norton (Original work published 1945).
Piaget, J. (1971a). In B. Walsh (Ed.), *Biology and knowledge*. Chicago: University of Chicago Press (Original work published 1967).
Piaget, J. (1971b). In E. Duckworth (Ed.), *Genetic epistemology*. New York: W.W. Norton (Original work published 1970).
Piaget, J. (1977). In M. Cook (Ed.), *The origin of intelligence in the child*. New York: Penguin (Original work published 1936).
Piaget, J. (2001). In R. L. Campbell (Ed.), *Studies in reflecting abstraction*. Hove, UK: Psychology Press (Original work published 1977).
Robbins, J. (1991). Echoes of 1939 on Capitol Hill. *Wall Street Journal*, *A12*.
Rosch, E. (1978). Principles of categorization. In E. Rosch, & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Erlbaum.
Rothbart, D. (1984). The semantics of metaphor and the structure of science. *Philosophy of Science*, *51*, 595–615.
Schön, D. A. (1963). *Displacement of concepts*. New York: Humanities Press.
Schön, D. A. (1979). Generative metaphor: A perspective on problem-setting in social policy. In A. Ortony (Ed.), *Metaphor and thought* (pp. 154–283). Cambridge: Cambridge University Press.
Solomon, M. (1992). Scientific rationality and human reasoning. *Philosophy of Science*, *59*, 439–455.

Toulmin, S. E. (1964). *The uses of argument*. Cambridge: Cambridge University Press.
Turbayne, C. M. (1962). *The myth of metaphor* (rev. ed.). Columbia: University of South Carolina Press.
von Wright, G. H. (1965). *The logical problem of induction* (2nd rev. ed.). Oxford: Oxford University Press.
Wegener, A. (1966). In J. Biram (Ed.), *The origin of continents and oceans* (4th ed). London: Dover (Original work published 1915).
Weitzenfeld, J. S. (1984). Valid reasoning by analogy. *Philosophy of Science*, *51*, 137–149.