

HOT COGNITION: MECHANISMS FOR MOTIVATED INFERENCE

*Paul Thagard and Ziva Kunda,
Psychology Department,
Princeton University,
Princeton, NJ 08544.*

Keywords: motivation, inference, computation, generalization.

Abstract

We present an implemented computational theory of motivated inference intended to account for a variety of experimental results. People make motivated inferences when their conclusions are biased by their general motives or goals. Our theory postulates four elements to account for such biasing. (1) A representation of the self, including attributes and motives. (2) A mechanism for evaluating the relevance of a potential conclusion to the motives of the self. (3) Mechanisms for motivated memory search to retrieve desired conceptions of the self and evidence supporting desired conclusions. (4) Inference rules with parameters that can be adjusted to encourage desired inferences and impede undesired ones.

1. INTRODUCTION

Pascal (1666) said that the heart has its reasons that reason does not know. But relatively little attention in cognitive science has been paid to "hot" cognition involving motivation and affect, as opposed to "cold" cognition involving problem solving, learning, and so on. Phenomena of motivated inference have been discussed by philosophers such as Fingarette (1969) and Haight (1980), and have been investigated experimentally by social psychologists (for reviews see Greenwald, 1980; Wicklund & Brehm, 1976). But aside from the path-breaking study by Abelson (1963), there has been little investigation of *how* motivated inference takes place, of the mechanisms by which motivations of the self influence the conclusions that it reaches (cf. Hastie, 1983; Sorrentino & Higgins, 1986).

Here we propose a computational model of motivated inference that accounts for a variety of phenomena that have been investigated empirically. Our account builds on the PI model of cold cognition developed by Thagard and Holyoak (1985; Holland, Holyoak, Nisbett, and Thagard, 1986). PI is a computational model of problem solving and learning.¹ We will describe an extension of PI, Motiv-PI, in which inferences such as generalization can be biased by the motivations of the system.

2. ELEMENTS OF A THEORY OF MOTIVATED INFERENCE

2.1. Kinds of Motivated Inference.

Our model is designed to provide an integrated account of several kinds of motivated inference.

(a) *Motivated changes of self-conceptions.* How people see themselves may be influenced by how they would like to see themselves. Thus people led to believe that extraversion is predictive of academic success come to view themselves as more extraverted (Kunda & Santoso, 1986).

(b) *Motivated changes of theories about the world.* People tend to generate those theories about the causal determinants of events that are most likely to support their goals. Thus people tend to believe that their own attributes are more predictive of happy marriage than are other people's attributes (Kunda, 1987). This allows them to

¹"PI" is short for "processes of induction" and is pronounced "pie".

maintain the belief that they will achieve a happy marriage.

Our model views both these phenomena as resulting from selective memory search among the wide array of relevant beliefs. Thus motivation helps to determine which self-conceptions will be accessed and which beliefs and what evidence pertaining to causal theories will be accessed.

(c) *Motivated changes of inferential rules.* Motivation affects the evaluation of evidence, so that individuals threatened by some evidence are less likely to believe it. Thus women who are heavy coffee drinkers were found to be particularly reluctant to believe that caffeine causes disease (Kunda, 1987). This reluctance may have been due to the application of particularly stringent inferential rules to the evaluation of the evidence, although to date there is no direct support for this notion. Our model assumes that motivation affects people's willingness to generalize by influencing the threshold required for generalization. Thus larger samples may be required to support generalizations that clash with one's goals.

(d) *Motivated changes of goals.* There is some evidence that when people realize that they are unlikely to obtain their goals, they diminish the importance of these goals to the self. Thus when individuals are outperformed by others on a given task, they come to consider that task as less important to their views of themselves (Tesser & Campbell, 1983). This allows them to maintain positive self-evaluation.

Thus motivation may affect inference by guiding the search among a wide array of potentially relevant beliefs about the world, other people, and the self, and by guiding the application of inferential rules.

2.2. Mechanisms.

We postulate four elements to describe the mechanisms underlying such motivated inferences:

1. *A representation of the self.* This should include *motives* of the self such as staying healthy and *attributes* of the self such as drinking coffee.

2. *A mechanism for evaluating the relevance of a potential conclusion to the motives of the self.* This will be an inference engine for tracing out the consequences of a potential conclusion and determining whether these have any impact on the motives of the self.

3. *Mechanisms for motivated memory search.* We hypothesize that motivation affects how people retrieve memories and make inferences about their own characteristics and goals and about evidence for and against potential inductive conclusions.

4. *Mechanisms for adjusting the parameters of inference rules.* These will distort the normal inference rules to ensure that inferences favorable to the self are more likely to be made and that unfavorable inferences are less likely to be made.

The interactions of these elements is depicted in figure 1. To begin, inference of a potential conclusion that would result from the application of some inference rule is triggered. The key question is: does the inference rule license the conclusion? In motivated inference this is not simply a matter of seeing whether the available evidence and the inference rule warrant the conclusion, because both the activation of evidence and the parameters of the rule may be influenced by motivation. First the potential relevance of the conclusion to the self is checked. If the conclusion is irrelevant to the self, motivation plays no role in inference and the standard cold inference rule is applied. But if the conclusion has positive consequences for the motivations of the self, then application of the inference rule is enhanced to make it more likely that the inference gets made. Motivated memory search encourages the retrieval of information that will support the conclusion, and evidence thresholds for applying the inference rule are lowered. If, on the other hand, the potential conclusion has negative consequences for the motivations of the self, then application of the rule will be impeded. Sometimes, however, evidence will be so overwhelming that inference will go through anyway, which prompts re-evaluation of the importance to the self of the implicated goals.

3. A COMPUTATIONAL MODEL

The elements just described provide only a sketch of a theory of motivated inference. We need to know a lot more about the structures and processes postulated in Figure 1. The best means currently available for specifying such structures and processes is to develop a detailed computer program that has data structures corresponding to the representations postulated and procedures corresponding to the processes postulated. For a theory of motivated inference, we need an account of the self and of the mechanisms of relevance evaluation, memory

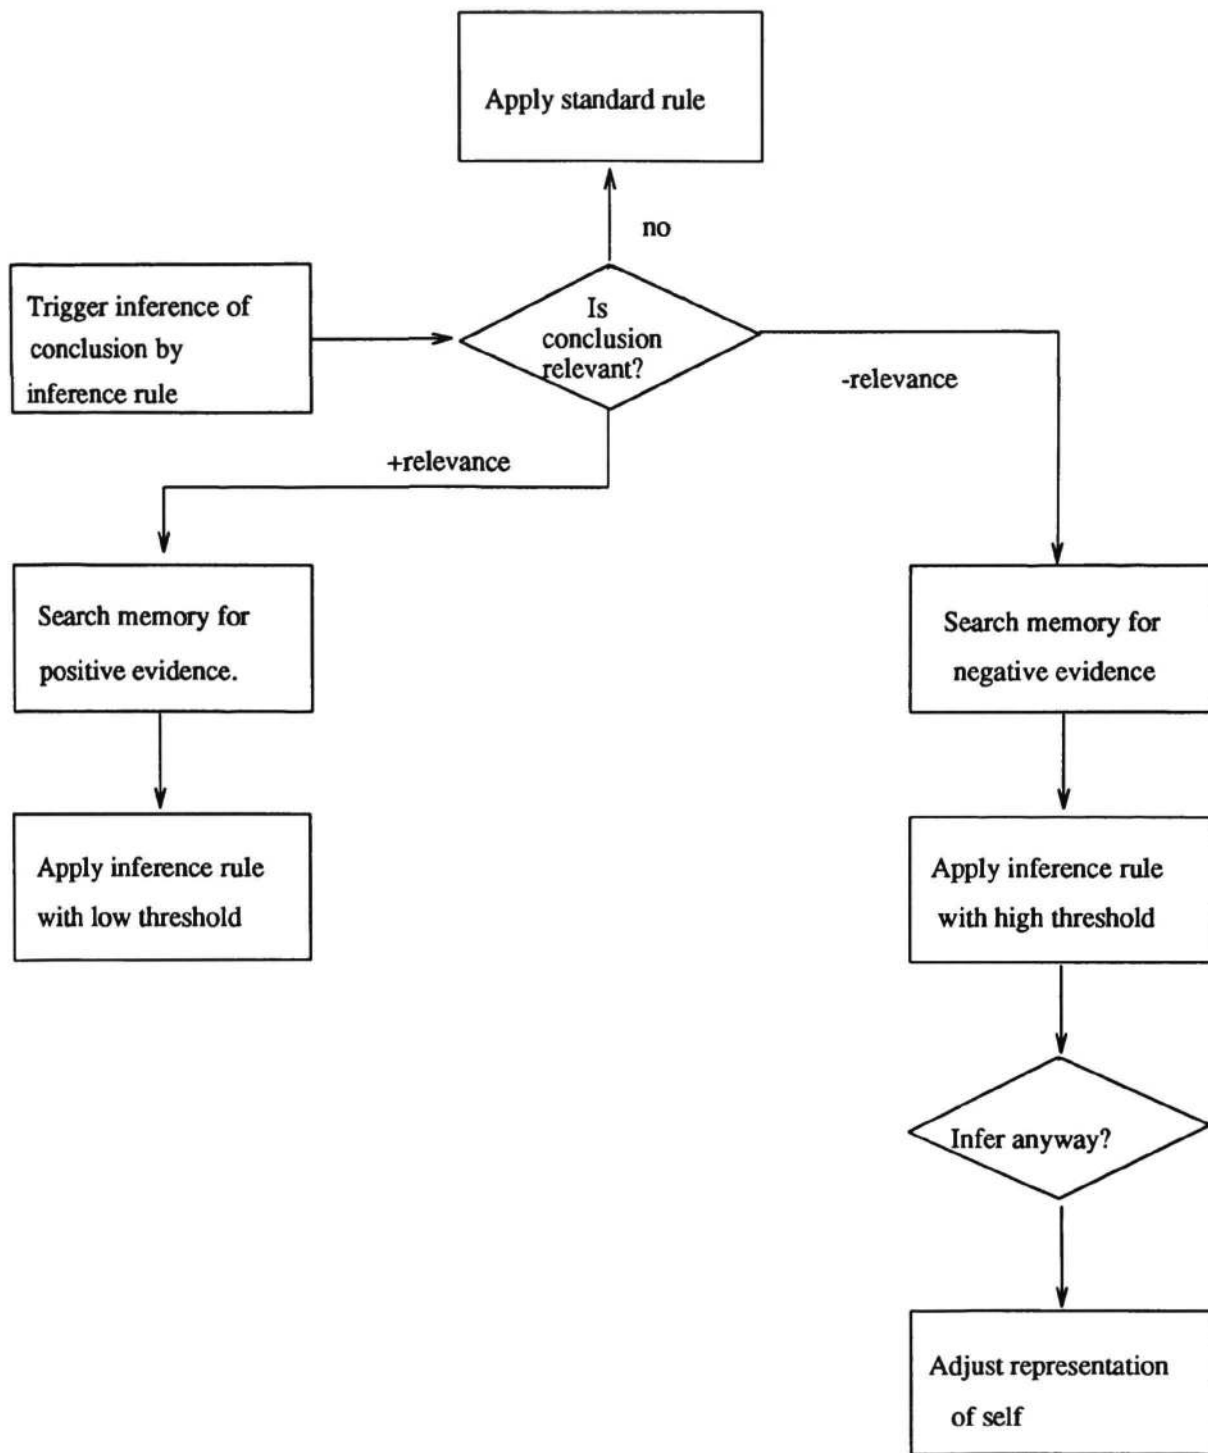


Figure 1: General model of motivated inference.

search, and distorted application of inference rules.

We now propose a computational model that builds on the PI model of problem solving and learning, a more detailed description of which can be found elsewhere (Holland, Holyoak, Nisbett, and Thagard, ch. 4). In PI, problem solving is a process of firing of rules and spreading activation of concepts. During problem solving, the current state of activation triggers various kinds of inductive inference, including generalization, abduction (inference to explanatory hypotheses), and concept formation.

3.1. The Self.

To expand PI into Motiv-PI, our model of motivated inference, we added a representation of the self to PI's set of knowledge structures which currently includes production rules, concepts (schemas), and messages (facts, propositions). Representation of the self has two crucial components: motives and attributes. The attributes of a self can include properties such as being a coffee drinker, a student, or a female. Its motives are its highest level goals, such as being healthy and happy. Figure 2 contains a simplified depiction of a self used in Motiv-PI's simulations. The activation of motives and attributes is variable, to allow differences in the degree to which they matter to the self at different times. For example, your motive to eat may be more active at one time than at another when you have just eaten, and your view of yourself as extraverted might be more active when you are at a party than when you are alone at your desk. Independent of activation, we postulate that motives have different priorities, so that reflection would tell you that being healthy is more important to you than being successful, even if at a particular moment you are overworking yourself to accomplish some career-related goal.

Self: Sandra

Attributes:	Motives:
Drinks_coffee:	Healthy:
Importance: 0.6	Priority: 0.9
Activation: 0.4	Activation: 0.4
Female:	Rich:
Importance: 0.5	Priority: 0.7
Activation: 0.6	Activation: 0.8
Student:	Happy:
Importance: 0.1	Priority: 0.9
Activation: 0.6	Activation: 0.5
Smokes:	Successful_career:
Importance: 0.7	Priority: 0.6
Activation: 0.4	Activation: 0.7
Had_non_working_mother:	
Importance: 0.1	
Activation: 0.4	
Had_close_father:	
Importance: 0.1	
Activation: 0.6	

Figure 2: A sample self.

Induction in Motiv-PI is triggered just as in PI, by the current state of activation of the problem-solving system. For example, the attempt to generalize that all A are B will be triggered by having the concepts of A and B simultaneously active as well as the information that something is both A and B. Providing examples of people who drink coffee and then develop fibrocystic disease would be sufficient to trigger the inference that coffee drinkers get the disease.

3.2. Determining Relevance of a Conclusion.

The second element of the theory of motivated inference mentioned in the last section is a means of determining the relevance of a potential conclusion to the self's motives. Motiv-PI does this by using the problem solving apparatus already present in PI. Determining relevance is a special case of problem solving, where the starting conditions of the problem consist of descriptions of the attributes of the self along with the potential conclusion, and the goals of the conclusion are the motives of the self. Normally, the point of attempting to solve a problem is to accomplish all the goals of the problem, but for determining relevance of a conclusion we want only to know which goals, i.e. which motives, have been accomplished. For example, to determine the relevance of the potential conclusion that Sandra has fibrocystic disease to the self Sandra, we set up a new problem whose starting conditions include the hypothetical proposition that she has the disease and whose goals are her motives:

Problem:

Start: (has_fibrocystic_disease (Sandra) projected_true)
... plus other attributes of Sandra.
Goals: (is_healthy (Sandra) true)
(is_happy (Sandra) true)
... plus other motives of Sandra.

Solving the problem consists of seeing whether any of Sandra's motives could be affected by the hypothetical start. In the case above, the conclusion that Sandra had fibrocystic disease would lead to the consequence that she is not healthy, indicating a negative evaluation of the conclusion. Figure 3 provides an overview of this process. In many cases, a chain of deductions will be necessary to calculate relevance. For example, to assess the motivational relevance of doing well in graduate school, Motiv-PI infers that doing well in graduate school will lead to a good first job, and getting a good first job will help to lead to a successful career, so that doing well in graduate school is positively relevant to having a successful career. Another evaluation, this time of having a stable marriage, is based on inferences that a stable marriage leads one to feel more secure and therefore to be happier. Thus evaluating the relevance of a possible conclusion depends on being able to trace out its consequences. We view it as an advantage of our model that no special mechanism is postulated for doing this, since the normal process of problem solving is used. Just as PI solves problems by simulating the effects of possible actions, Motiv-PI determines relevance to the self by using information stored in concepts and rules to infer the consequences for the self a potential conclusion. The problem solver notes what motives are accomplished by a potential conclusion, yielding a numerical total that takes into account both the priority of the different goals accomplished and their degree of activation.

3.3. Motivated Memory Search.

People possess a broad array of different and sometimes contradictory beliefs about themselves, others, and the world. Motivation can determine which of these beliefs they will retrieve so as to support or hinder potential inductive conclusions. Motiv-PI implements such motivated use of memory very naturally because of the subgoal mechanism in PI. PI's basic problem solving mechanism is forward chaining, matching production rules such as *If x is sociable, then x is extraverted* against messages such as *Sandra is sociable* to generate conclusions such as *Sandra is extraverted*. But if the system is motivated to show that Sandra is extraverted, it sets that as a goal and chains backwards to activate information about her being sociable. Further subgoaling occurs using an active rule that says that people who go to lots of parties are sociable, leading the system to ask if Sandra goes to lots of parties. Retrieving this information will then make possible the inference that Sandra is sociable and thus is extraverted. If the desired conclusion is reached, the search stops, so the system does not go on to find contradictory information that might imply that the self is introverted. This contradictory information will therefore not be accessed unless it is available to begin with or is activated as a side-effect of the main search. If instead the system were motivated to show that Sandra is introverted, it could do so by activating different rules, such as that shy people are introverted, and different facts, such as episodes where in fact Sandra was shy. Thus motivated memory retrieval occurs using the subgoaling which is an integral part of the problem solving operation of PI.

The same mechanisms govern motivated retrieval of evidence. In motivated generalization, you can be motivated to believe that all A's are B's, in which case you want to retrieve as many A's that are B's as possible. Or, if you are motivated not to form this generalization, then you will want to find examples of A's that are not B's. Motivated retrieval in the former case consists of giving PI's problem solver the goals of finding things that

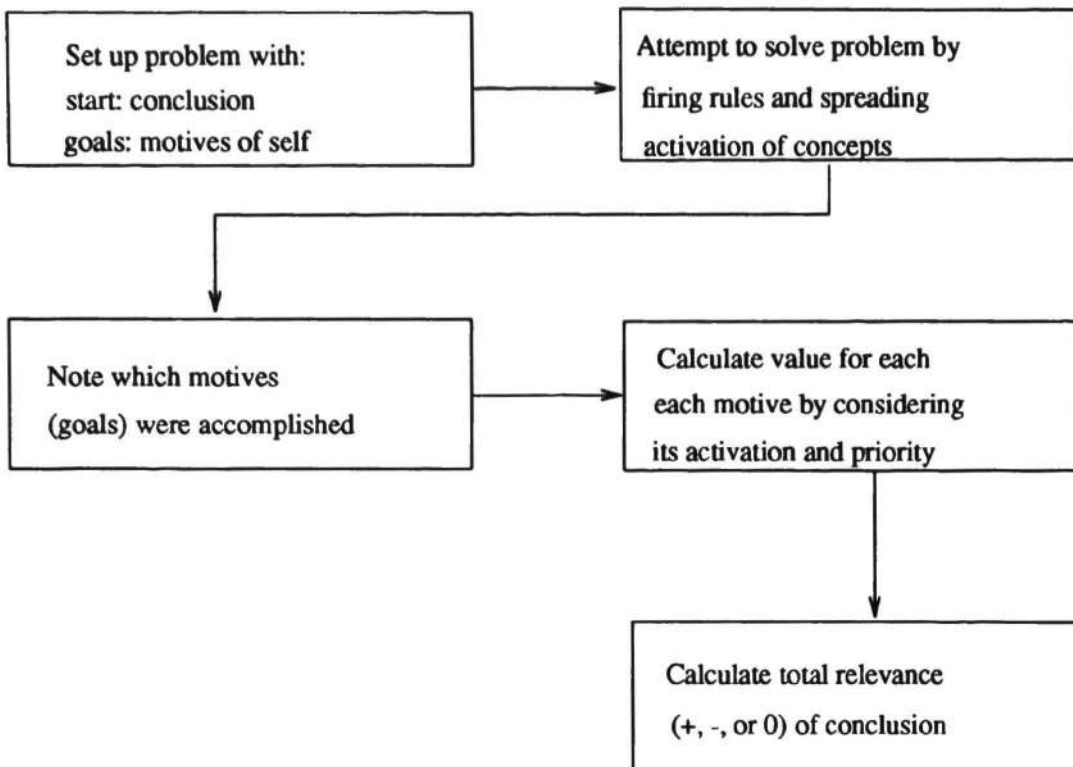


Figure 3: Calculating the relevance of a potential conclusion.

are A's and B's and storing this information away with the the concepts A and B for use as evidence in favor of the generalization. For example, if you want to infer that extraverts are successful, you will do a memory search to construct and retrieve examples of successful extraverts, which may involve the realization that some people whom you had not previously encoded in these terms do in fact fall under both categories.

In Motiv-PI, the extent of the motivated search for evidence is a function of how motivated the self is to form a particular conclusion. The amount of inferences PI's problem solver will employ to try to turn up desired information is a function of the extent to which the potential conclusion is relevant to the self, taking into account both the importance and the degree of activation of the relevant motives. Experimental results suggest that high levels of motivation are required for motivated inference to occur (Kunda, 1987).

3.4. Motivated Generalization.

The calculated motivational relevance of a potential conclusion and motivated memory search can be used by Motiv-PI to distort various inference processes. Here we will concentrate on generalization, since the kinds of inferences studied experimentally by Kunda fall most appropriately under that heading. In PI, generalization is done in accord with a theory developed by Thagard and Nisbett (1982) and tested experimentally by Nisbett, Krantz, Jepson and Kunda (1983). People's willingness to generalize from instances is a function both of the number of instances that provide evidence for the rule to be formed and of background knowledge about variability. You are more prone to generalize that all instances of a new kind of metal burn with a green flame than you are to generalize that all instances of a new kind of bird are blue, since you know that birds are more variable with respect to color than metals are with respect to combustion properties. PI generalizes that all A's are B only when it calculates that a combined measure of the number of instances of A's that are B's and the invariability of A's with respect to B's exceeds a given threshold. (For further discussion of generalization and variability, see Holland, Holyoak, Nisbett, and Thagard, 1986, ch. 8.)

Motiv-PI uses the same considerations when attempting to form generalizations, but motives influence the process in the ways summarized in Figure 4 and described below. A rule "All A are B" is typically relevant to a self if (1) the self has attribute A, and (2) B is motivationally relevant to the self -- the self desires or fears being B. For example, the rule "Women who drink coffee get fibrocystic disease" is only relevant to someone who is female, drinks coffee, and cares about being healthy.

When attempting to form the generalization "all A are B" Motiv-PI first determines whether being B would be relevant to the self. If there is no relevance, then it does generalization just as in PI, with the standard threshold for number of instances and variability. If being B is positively relevant to the self, leading to satisfaction of its motives, then Motiv-PI first attempts to show that the self is A, using the motivated memory search described above. For example, when considering the potential generalization by Sandra that all extraverts are successful, it first determines that being successful is positively relevant to Sandra, and then tries to show that Sandra is extraverted. Since PI simulates the parallelism of numerous rules firing and subgoaling at once, the following subgoaling chains occur simultaneously:

extravert <- friendly <- has many friends (Sandra)

extravert <- sociable <- goes to parties (Sandra)

extravert <- outgoing <- talks to strangers (Sandra)

If B is relevant to the self and the search just described shows that the self is A, then the system fosters generalization that all A's are B's in two ways. First, it attempts to find as many examples of A's that are B's as possible, in order to get above the threshold for the number of instances required for generalization at a given level of variability. In the extravert/success example, the system searches for instances of individuals who are both extraverted and successful by setting itself the subgoals of finding things that fall under these categories. Second, the system adjusts the threshold of the inference rule used in generalization to require fewer instances given the calculated variability.

If being B is a property such as getting breast disease that has negative consequences for the self's motives, then Motiv-PI tries to show that the self is not A. For example, if generalization is triggered by examples of extraverts that failed, Motiv-PI attempts to show that the self is not extraverted. If the self is found nevertheless to be extraverted, then the system attempts in two ways to block generalization that extraverts fail. First, it does a memory search by subgoaling to try to find examples of non-failing extraverts. (PI's normal generalization mechanism blocks inference when such counterexamples are available.) Second, just as positive motivated generalization uses a lower threshold of number of instances and variability, in this negative case a higher threshold is used to impede generalization.

These mechanisms do not, however, guarantee that the undesired generalization will be blocked, since if sufficient evidence is found the generalization will be made nevertheless. In this case, Motiv-PI adjusts the representation of the self to make it less concerned with the bad consequences that derive from the conclusion that all A are B. This form of rationalization consists of deciding that perhaps being B is not so bad after all. In other words, Motive-PI reduces the priority of the motive that B affected. To continue the above example, after concluding that its personality will make it less likely to succeed, the system will decide that success is not all that important to it.

Currently, the variability calculation is not affected by motivation, but conceivably the memory search on which the calculation of variability is based could be substantially affected by motivation. Variability calculations depend on selection of appropriate reference classes. To generalize that all shreebles (a new kind of bird) are blue, PI calculates the background variability of birds with respect to color, but a more sophisticated program would pick reference classes less automatically. It might, for example, be appropriate to consider more specific classes, calculating the variability of tropical birds with respect to primary colors. Perhaps motivation plays a role in searching for reference classes to generate variability estimates that foster or impede generalization. More experimental studies are needed, however, to show whether motivation has an effect on variability calculations and instance retrieval.

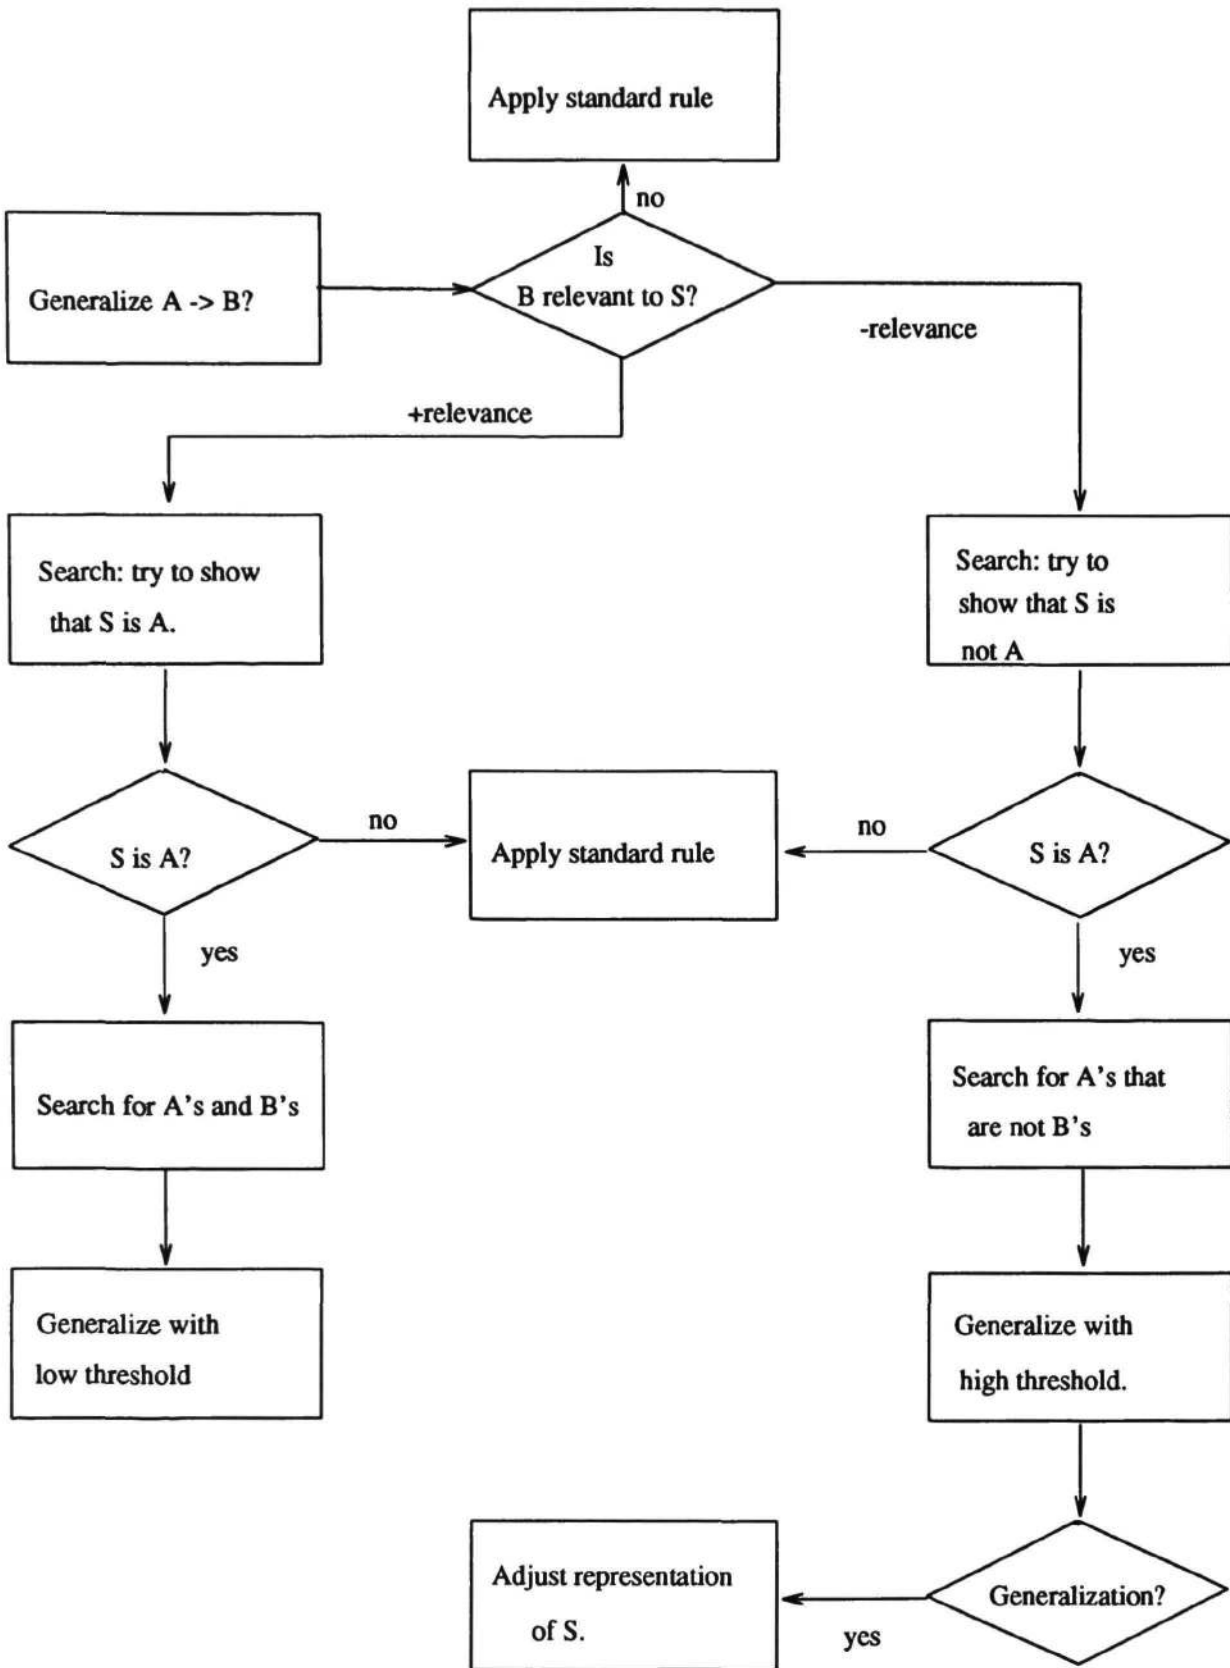


Figure 4: Motivated generalization by self S that all A's are B's.

3.5. Motivated Deduction and Hypothesis Evaluation

Other types of inference besides generalization may also be subject to motivational biases, although these have not yet been studied experimentally or implemented computationally. One might think that deduction is immune from motivational biases, since any conclusion that follows deductively from true premises has to be true. But in any realistic processing system, deduction has to be constrained pragmatically, since any system that made all possible deductive inferences would suffer a combinatorial explosion. In PI, what deductions get made during problem solving is a function of pragmatics such as the problem solving context, particularly the goals to be accomplished. Thus the memory search described above is a kind of motivated deduction.

PI also performs abduction, that is, formation of explanatory hypotheses. The simplest kind offers an explanation of a fact that something has a property B by hypothesizing, on the basis of a rule *All A are B* that it is A. For example, to explain why a friend is exhausted and dissipated, you might abduce that he or she is on drugs, since being on drugs can lead to dissipation. Evaluation of hypotheses formed by abduction is by inference to the best explanation, which PI performs by compiling lists of alternative hypotheses and facts to be explained, and evaluating which of the hypotheses most comprehensively and simply explains the evidence (Thagard, forthcoming, ch. 5). It is easy to see how motivation could distort inference to the best explanation: if you are motivated to accept a hypothesis, then you might be less thorough in searching for alternative hypotheses and more prepared to accept the hypothesis on less evidence. Indeed, it has been shown that the success of a liked person is attributed to the person's ability, whereas the success of a disliked person is attributed to the ease of the task (Regan, Strauss, & Fazio, 1974). For Motiv-PI, alterations could easily be made to bias inference to the best explanation.

3.6. Limitations of Motiv-PI.

Motiv-PI, including representation of the self, calculation of motivational consequences, motivated generalization, and rationalization is implemented in Common LISP and runs in conjunction with PI. A fuller implementation would add features such as the following:

1. A dynamic self, in which the degree of activation of its attributes and motives is updated by the program itself.
2. Motivated variability judgments.
3. Motivated inference to the best explanation.
4. More sophisticated rationalization by adjusting motive importance indirectly through motivated memory search: instead of merely reducing the priority of a motive, the system would try to retrieve evidence that would diminish the inferred importance of the motive.
5. Greater parallelism so that different motivated memory searches could be simulated as occurring at the same time.

4. COMPARISON WITH OTHER VIEWS

4.1. Computational Models of Motivation and Affect

Abelson (1963) proposed an interesting model of sentence evaluation that had some of the features we have discussed in connection with motivated inference. He described a system for cognitive balancing, in which a sentence that enters thought is evaluated and unbalanced sentences are subject to further processing before they are stored. For example, the sentence "My good friend is a murderer" is unbalanced because its subject and predicate get very different evaluations. Here rationalization consists of modifying the subject or predicate in some way to restore balance. Our account can be thought of as a way of achieving balance, or reducing dissonance, between the motives of the self and the inferences it makes, but the empirical and computational work described here are novel in that they shift the focus to the memory and inference processes through which balance is maintained.

We can only briefly mention some other computational studies of affect. Wegman (1985) offers a computational model of Freud's theories of abreaction and repression. Ortony (1986) has proposed a model of event evaluation which is somewhat similar to our account of evaluating the relevance of potential conclusions concerning the self. Dyer (1983) and Hovy (1986) have discussed the relevance of affect in text processing.

4.2. Philosophical Applications

Philosophers have paid much attention to phenomena closely related to motivated inference: weakness of will and self-deception (Davidson 1980, Fingarette 1969, Haight 1980, Martin 1985). One central question has been whether these phenomena are possible at all. On a simple view of the mind, it can be difficult to understand how an agent could act other than in its own interests or could deceive itself. However, the structures and processes used in Motiv-PI make it easy to see how weakness of will and self-deception can at least be possible. It is another question, that ought to be answered experimentally, whether people are actually subject to them.

We see self-deception as an extreme case of motivated inference leading to inconsistency. In self-deception, one makes an inference that at some level one knows to be false. If I infer that my finances are solid even though I know that bankruptcy is imminent, then I am guilty of deceiving myself. On some philosophical views, it is hard to see how such a contradiction could exist. But it is perfectly consistent with PI's representations and processes that a system be inconsistent without realizing it. Only a portion of the system's beliefs need be active at any time, so it is very possible to infer a belief that is active but inconsistent with another belief that is not active. According to Audi (1985), in typical cases of self-deception we have not only a motivated inference of a belief that contradicts an existing one, but also removal from consciousness of the original belief, which sounds a lot like Freudian repression. A repression mechanism could easily be added to Motiv-PI. The system could evaluate the relevance to the self of each active message (proposition) and de-activate the negative ones, but this seems much too strong: there is no reason to expect that all information will be active at a given time, so self-deception would seem to be possible without a repression mechanism. Motivated memory retrieval and inference suffice to produce the phenomenon.

On our view, weakness of will can be understood in terms of the difference between priority and activation of motives. Weakness of will occurs when a decision is motivated by desires that are more active than ones to which reflection gives a higher priority. To take an extreme case, consider a cocaine addict fighting a craving for the drug. For physiological reasons, the motive to get some cocaine is far more active than long-term motives such as being healthy, happy, and successful, so that even though the addict would on reflection give higher priority to the latter motives, the desire for cocaine wins out.

5. CONCLUSION

We have described mechanisms for motivated inference that have been worked out in enough detail to be implemented computationally. We close by sketching a practical motive for these investigations. Motivated inference may sometimes be harmless, or even beneficial, despite fostering false conclusions. For example, believing on little evidence that you will succeed at some difficult task may well improve your performance. But motivated inferences can also be dangerous, even life threatening. The smoker who fails to conclude that smoking contributes to cancer may pay for this by incurring the disease. We hope that the investigation of the processes underlying motivated inference will eventually lead to the development of methods for helping people avoid hazardous reasonings from the heart.

6. REFERENCES

- Abelson, R. (1963). Computer simulation of "hot" cognition. In S. Tomkins (Ed.) *Computer Simulation of Personality*. New York: John Wiley and Sons. 277-298.
- Audi, R. (1985). Self-deception and rationality. In M. Martin (Ed.) *Self-deception and self-understanding*. Lawrence, KS: University Press of Kansas. 169-194.
- Davidson, D. (1980). *Essays on action and events*. Oxford: Oxford University Press.
- Dyer, M. (1983). The role of affect in narratives. *Cognitive Science*, 7, 211-242.
- Fingarette, H. (1969). *Self-deception*. London: Routledge and Kegan Paul.
- Greenwald, A. A. (1980). The totalitarian ego: Fabrication and revision of personal history. *American Psychologist*, 35, 603-618.
- Haight, M. (1980). *A study of self-deception*. Brighton, Sussex: The Harvester Press.
- Hastie, R. (1983). Social inference. *Annual Review of Psychology*, 34, 511-542.
- Holland, J., Holyoak, K., Nisbett, R., and Thagard, P. (1986). *Induction: Processes of inference, learning, and discovery*. Cambridge, MA: Bradford Books/MIT Press.

- Hovy, E. (1986). Putting affect into text. *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum. 669-675.
- Kunda, Z. (1987). Motivation and inference: Self-serving generation and evaluation of causal theories. *Journal of Personality and Social Psychology*. In press.
- Kunda, Z. & Santioso, R. (1986). Motivated changes in the self-concept. Unpublished manuscript, Princeton University.
- Martin, M. (Ed.). *Self-deception and self-understanding*. Lawrence, KS: University Press of Kansas.
- Nisbett, R. E., Krantz, D. H., Jepson, D., and Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, *90*, 339-363.
- Ortony, A. (1986). Unpublished manuscript. University of Illinois.
- Pascal, B. (1966). *Pensées*. Harmondsworth, Middlesex: Penguin.
- Regan, D. T., Strauss, E., & Fazio, R. (1974). Liking and the attribution process. *Journal of Experimental Social Psychology*, *10*, 385-392.
- Sorrentino, R. M. & Higgins, E. T. (1986). Motivation and cognition: Warming up to synergism. In R.M. Sorrentino and E.T. Higgins (Eds.), *Handbook of motivation and cognition*. New York: Guilford Press.
- Tesser, A., & Campbell, J. (1983). Self-definition and self-evaluation maintenance. In J. Suls & A. Greenwald (Eds.), *Social psychological perspectives on the self* (Vol. 2). Hillsdale, NJ: Erlbaum.
- Thagard, P., and Holyoak, K. (1985). Discovering the wave theory of sound: Induction in the context of problem solving. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*. Los Altos: Morgan Kaufmann. 610-612.
- Thagard, P., and Nisbett, R. E. (1982). Variability and confirmation. *Philosophical Studies*, *42*, 379-394.
- Thagard, P. (forthcoming). *Computational philosophy of science*. Cambridge, MA: Bradford Books/MIT Press.
- Wegman, C. (1985). *Psychoanalysis and cognitive theory*. Orlando, FL: Academic Press.
- Wicklund, R. A. & Brehm, J. W. (1976). *Perspectives on cognitive dissonance*. Hillsdale, NJ: Erlbaum.