

Belief Relativity

Donald Rose

Department of Information & Computer Science
University of California
Irvine CA 92717
drose@ics.uci.edu

Abstract

This paper describes a model of belief systems called *belief relativity* (BR), which addresses the relationships and structure of knowledge held by multiple interacting agents. This paradigm uses *belief reference frames* (b-frames) as the main unit of belief spaces, within which an agent's beliefs are stored. BR is concerned with how beliefs are created and revised, how they influence each other within or between b-frames, and how one searches for b-frames that are useful (e.g., that remove contradictions). BR also deals with degrees of belief, propagated along the influences that relate beliefs and b-frames to each other. BR attempts to combine the best features of these ideas into a unified, synergistic framework.

Introduction

This paper presents a model of belief systems called *belief relativity* (BR). The model consists of two main interacting elements. First is the notion of *belief reference frames* (b-frames), within which an agent's beliefs reside along with their *influences* – i.e., supporting beliefs, and any other more indirect evidence. This area of BR is concerned with how beliefs are created and revised, how they influence each other within or between b-frames, and how one searches for useful b-frames (e.g., that remove contradictions). This area has been studied in research on TMS-based systems (Doyle 1979) as well as systems related to ATMS (de Kleer 1984; Rose & Langley 1986). The second notion is *degrees of belief*, propagated along the aforementioned influences between beliefs. This area of BR is concerned with how one can utilize various types of rules to infer the degree to which a belief fits into a given property class. This has been addressed by work in fuzzy systems (Zadeh 1965), as well as other methods for mathematical computation of evidence (Shafer 1976). BR attempts to combine some of the best features of these existing belief models – as well as new ideas – into a unified, synergistic framework, one which can hopefully apply to a wider class of problems, or offer insights into how they may eventually be addressed.

Beliefs and Reference Frames

One of the basic premises of BR is that the reasoning of intelligent agents should be modelled as reasoning with several types of beliefs, all of which can be grouped into reference frames (b-frames). In BR, beliefs are generally represented as follows:

$\langle \text{property } P \rangle (\langle \text{b-frame} \rangle, \langle \text{assertion } A \rangle) = \langle \text{degree that } P \text{ holds for } A \rangle .$

The BR model assumes that an assertion is any unit of knowledge that can be attached to a property; for instance, agents can attach the property “relevance” to an assertion. Beliefs can in turn act as assertions to be acted upon by other beliefs, as discussed later. In such nested beliefs, all b-frames are that of the outermost belief, unless stated otherwise. BR allows agents to adopt many b-frames over time, depending on the state of one's reasoning (e.g., when exploring hypotheses within a temporary b-frame that can be deactivated at a later time).

The most basic frame concerns belief *influence* (i.e., the frame that indicates causes, or inferential support, or evidence). A link from a belief X to another belief Y indicates that X influences Y in some fashion; belief X may not directly lead to the inference of Y, yet it may still influence Y in some other indirect manner. The link from X to Y can carry a *degree* of influence, depending on the rules that are used by an agent to perform inferencing. Sets of beliefs, and links that represent the degree of influence, form networks, which have been proposed and analyzed elsewhere (*influence diagrams* (Howard & Matheson 1984) and *belief networks* (Pearl 1988)). For example, the b-frame of a nonsmoker NS might compute that smoking highly affects cancer rates – i.e., $\text{influence}(\text{NS}, (\text{smoking}, \text{cancer})) = 0.8$, whereas the b-frame of a smoker S might feel the influence is less – i.e., $\text{influence}(\text{S}, (\text{smoking}, \text{cancer})) = 0.4$.

Note that these examples of *causal influence beliefs* do not indicate inferential support, only causal support. The special case where influences have only 0 and 1 as possible strength values (i.e., degrees of influence are either all or none), and the influences indicate direct inferential support (i.e., can be viewed as inference histories) is analogous to a number of systems for automated reasoning – e.g., TMS, ATMS or REVOLVER (Rose 1989). Examples of *support influence beliefs* might be: $\text{influence}(\text{NS}, (\text{B1}, \text{B2})) = 0.7$, and $\text{influence}(\text{S}, (\text{B1}, \text{B2})) = 0.4$, where B1 and B2 are other beliefs. B1 might be $\text{influence}(\text{surgeon-general}, (\text{smoking}, \text{cancer})) = 0.9$ and B2 might be $\text{influence}(\text{NS}, (\text{smoking}, \text{cancer})) = 0.9$. In other words, NS's believed influence of smoking on cancer is *itself influenced by* (e.g., inferred from) the surgeon-general's belief in this smoking-cancer connection, and NS believes this support with strength 0.7. S does not hold this B1-B2 support as strongly (only 0.4); perhaps S feels the surgeon-general's belief only partly influences NS's stance on the smoking-cancer connection. Later evidence may prove S

right (e.g., influences that neither S nor NS immediately recognized, such as old experiential influences that over time have lost nearly all their relevance to more current beliefs). Also note that B1 and B2 need not be influence beliefs, but can also be other types, such as relevance; these new beliefs are discussed below. In general, the nesting of beliefs within beliefs, and the interplay between beliefs of multiple agents, are two of the strongest characteristics and advantages of the BR framework.

Although influence networks are useful tools, the framework proposed in this paper does not limit beliefs to influences. A second type of belief would concern the degree of *certainty* in the beliefs of the influence b-frame. For example, the smoker S above might feel *highly* certain about his belief in the *low* influence of smoking on cancer. This could be represented as:

$\text{certainty}(S, \text{influence}(S, (\text{smoking,cancer})) = 0.4) = 0.99.$

However, S might *also* feel highly certain about *how NS feels about the influence of smoking on cancer*; i.e.,

$\text{certainty}(S, \text{influence}(NS, (\text{smoking,cancer})) = 0.8) = 0.9.$

The third main belief type in our framework deals with *relevance*. As an example, NS, being a nonsmoker, might hardly ever think about smoking issues, and hence not find S's beliefs very relevant; e.g.,

$\text{relevance}(NS, (\text{influence}(S, (\text{smoking,cancer})) = 0.4) = 0.1.$ In fact, NS's own beliefs on the subject may also have little relevance to NS:

$\text{relevance}(NS, (\text{influence}(NS, (\text{smoking,cancer})) = 0.8) = 0.1.$

However, like the other belief types, the degree of relevance might change as new beliefs become part of one's current b-frame. A good example is the inclusion of *goal* beliefs. Goals can be used to increase the degree of relevance of other beliefs. For example, the need to debate a smoker might increase NS's relevance for his beliefs pertaining to smoking issues.

One of the interesting aspects of including relevance and certainty in the representation of agent's belief systems is that they are likely to be intricately *intertwined*. In general, the degree of belief in certain assertions may lead to inferences regarding their relevance, and the degree of relevance can help drive which degrees of belief get propagated and when. Let us look at additional examples which hopefully illustrate these ideas. First, note that relevance may rise if one's degree of certainty rises. For example, the more I believe I have a cold, the more relevant "cold" assertions may get (e.g., remedies; rules for getting well). Alternatively, relevance may decline if degree of belief rises. For example, the more a Mom believes her son is safe, the less she might worry about her son's safety. A third case could arise where, for instance, an agent Y finds agent X's low degree of belief (say in Y's abilities) a very relevant issue.

Note that relevance is a *relative* quality, depending on one's b-frame. As another example, an agent X's assertions (and degrees of certainty in their causal connections) may make a frame F very relevant to X, but agent Y's assertions and beliefs may make it irrelevant to Y. More specifically, suppose an influence frame consists of all beliefs and causal influence links regarding the Kennedy murder conspiracy theory T, and suppose an author A creates, updates, and strongly believes T. T's assertions are thus highly relevant to A, and are almost always active. A kid K might know

of T but not believe it much; hence T's assertions have low relevance to K, and are almost always inactive. Lastly, a hippie H might believe T's assertions as highly as A, but have no vested interest in T. Hence T's assertions have low relevance to H, and are almost always inactive.

Having different belief types may also allow more accurate modelling of cognitive phenomena. For example, non-increasing relevance in one part of the influence b-frame's assertions may gradually lead to the virtual deletion/forgetting of these assertions. That is, it may become last on the list of things to think about. This behavior might prove quite valuable. For instance, if one is dealing in limited-memory scenarios (e.g., an autonomous long-term learning system without hardware expandability, such as in space domains), one could implement a Least-Recently-Used-type strategy; this method would explicitly delete those assertions that are either least relevant or least believed, taking into account in how long such a status held.

In summary, utilizing these and other types of interacting belief types should provide a more powerful framework for modelling complex belief reasoning phenomena. By modelling not only how highly one believes an influence, but also the relevance of that influence to one or more agents, the BR model imposes a useful structure on an agent's belief space, which would hopefully enable improved inferring and revision performance.

Propagating Degrees of Belief

In a general implementation of the BR model, any belief could be used as input to the degree-of-belief rule of any property (e.g., attractiveness; tallness; relevance). For example, one can have a rule relating the weight of someone (input) to the certainty that that weight is a member of the property class *heavy* (output). (Note, however, that relevance is currently the only property that can apply to any belief.) These "degree-of-property" beliefs can then be described by other properties, to any level of desired nesting. An advantage of BR is that degree-of-property inference rules are independent of the influences that determine which beliefs apply to each other, and can be of any type (e.g., fuzzy/probabilistic membership functions; qualitative ranges; binary decision rules).

For example, a qualitative rule might map beliefs to qualitative values (e.g., if $\text{height}(X) \geq 6'4"$, then $\text{tallness} = \text{very-tall}$; else if $\text{height}(X) \geq 5'8"$, then $\text{tall} = \text{tall}$; else $\text{tallness} = \text{short}$). Fuzzy membership functions might map beliefs to a point on along a continuous curve (e.g., $\text{tallness}(\text{height}(X) = 6'4") = 0.9$). Note that the belief being input *can itself represent a degree of belief*. That is, every application of a belief strength rule by a b-frame results in a new belief, which can potentially become a new input to any other b-frame (e.g., a new domain value in some other membership function). For instance, an agent X might not find the tallness of anybody that relevant; e.g.,

$\text{relevance}(X, \text{tallness}(X, \text{height}(Z) = 6'4")) = 0.9) = 0.2)$.

In general, the same belief can be processed by *different b-frames* (i.e., different agents, or different hypothetical views by the same user), and these different b-frames may employ different rules for determining degree-of-belief for different properties. Thus, if X worked for the Guinness Book of World Records, and Y told X about an extremely tall person, X might feel this quite relevant; e.g.:

$\text{relevance}(X, \text{tallness}(Y, \text{height}(Z) = ?)) = 0.99999) = 0.9$).

Note that it is Y's tallness belief that excites X, not Z's actual height; X's relevance belief represents the view from X's current b-frame, where exact knowledge of Z's numerical height is not known yet (only exists in Y's b-frame). Once X sees Z directly ($\text{height}(Z) = 6'6''$), X still thinks Z is tall (e.g., $\text{tallness}(X, \text{height}(Z) = 6'6'') = 0.9$), but this is a lesser degree of tallness than Y believes. Now using her own tallness judgement – the tallness belief about Z from her own b-frame – X revises her excitement downward; e.g.:

$\text{relevance}(X, \text{tallness}(X, \text{height}(Z) = 6'6'') = 0.9) = 0.3$).

Let us now look at a more complex example, where agents Jack and Kelly each use a b-frame (call these J and K). K has two fuzzy membership functions: one for determining her certainty in the heaviness of someone, another for determining the relevance of heaviness to her. We will assume J has only one membership function, a fuzzy view of K's heaviness-relevance belief about him. That is, J's relevance function determines the relevance of how strongly K finds J's weight relevant to her. A chain of influences has been constructed – from J's weight (belief W), to K's belief in W's degree of heaviness (belief X), to K's belief that X is relevant to her (belief Y), to J's belief that Y is relevant to him (belief Z).

The resultant computations of degrees of belief for W, X, Y and Z can be summarized as follows. If J is 190 pounds, let us suppose K's function finds him quite heavy (level 0.9). This level then becomes the basis to compute the relevance of this heaviness level to her; let us suppose her relevance function indicates her attraction to men over 170 pounds decreases fast; at 190 her attraction level is 0.3 (hence her degree of attraction to J is 0.3). Finally, this degree of attraction that K feels for J is input for J's function for determining his attraction for K. Since K feels attraction for J of only 0.3, J feels attraction for K (assuming a linear function) at the same low level (0.3).

If we were modelling height instead of weight, the influences would remain static over time, since height is not changeable. However, modelling weight allows changes to be made and propagated through the membership functions. If J's weight increases, K increases her certainty in J's heaviness. This belief then causes a decrease in her attraction to J. Finally, this lower degree of attraction that K feels for J causes J to decrease his attraction for K.

Note that J's choice of using an attraction-for-K rule that uses K's attraction-for-J as its sole influence sets J up for the effects of *cognitive dissonance* (Festinger 1957). That is, J essentially likes K if she likes him back, but once her interest falls, his does too. The richness of allowing multiple interacting b-frames, with interacting degree-of-belief rules, at user-controlled levels of nesting, will hopefully allow agents to model (and, if desired, perform) many qualitative classes of cognitive behavior. An even harder problem is how an agent might decide a belief propagation rule is yielding nonoptimal results, then begin a search for different influences (i.e., change b-frames, and hence change belief propagation rules). The ideas in this paper should provide a starting point from which to build such capabilities into future systems.

Revision in B-Frames

Several methods currently exist for doing belief revision on influence networks, such as in representations where premise beliefs support other inferred beliefs (e.g., TMS, ATMS, REVOLVER). However, these revision methods tackle only one part of the b-frame proposed above: reasoning within one b-frame, and doing so only on influence beliefs and links. A more general method of belief revision would enable operations that involve multiple b-frames (e.g., how to change from one to another, and when), as well as certainty and relevance beliefs and their associated links. Possible revision operators might be: add or delete a belief (essentially the operators of REVOLVER); modify a belief property (e.g., increase or decrease certainty or relevance); change b-frames (e.g., adopt a larger or smaller set of beliefs and links).

One issue involves the triggering of a *change of frames*. For example, an agent might deem a particular degree of belief to be too neutral (i.e., at or near 0.5). The agent may then decide to form an opinion about this link, one way or the other. Similarly, an agent may want to form an opinion based on beliefs (or degrees of belief) held by other observers; that is, one might analyze beliefs from other agents' b-frames to influence the beliefs or links of one's own b-frame.

To accomplish these revision goals, one might adopt a larger (wider scope) b-frame, in the hope of obtaining more evidence about new beliefs and using this knowledge to revise the beliefs (or belief properties) of one's old b-frame. One method for revising one's b-frame to a larger belief set might be to increase the relevance of other beliefs – in this case, those observed to be present in other agent's b-frames. One could begin by increasing the relevance of those external beliefs whose properties have the highest values in the other agents' b-frames. This embodies the intuitive heuristic: "pay more attention to the strongest outside opinions" (all other factors being equal).

In general, one might view belief revision as a *search through the space of b-frames*, using the operators proposed earlier. In most domains, there should also exist *constraints* one can use to help control this search. For example, one might utilize constraints on what evidence is needed to resolve a current conflict. The need to resolve a conflict between beliefs X and Y should lead an agent to limit revision search to frames involving both X and Y; the goal b-frame is one where the X-Y conflict is resolved. A more complex scenario would be to perform the above search, but do it so that another agent's conflicts are resolved; i.e., put oneself into another's head to help resolve a conflict. There are several states that changing one's b-frame can result in. In general, new *types* of beliefs, new *degrees* of belief, or both, can be created or inferred due to a change in b-frame.

Note that allowing degrees of belief on links means a reworking of the concept of contradiction. In the special case where all degrees of belief are either zero or one, a belief B is either held or not; holding both B and not(B) would be a clear contradiction. However, contradictions are more "fuzzy" if beliefs are held in degrees. For example, one might hold B with certainty 0.3 and not(B) with certainty 0.4; should a pair of beliefs with these certainty levels be deemed a contradiction?

To address this, an interesting revision strategy might be to use the degree of difference between two beliefs' certainty levels to influence the degree of relevance of this difference.

That is, as the difference in certainties between two conflicting assertions becomes closer to the standard definition of contradiction (holding B and not(B) with certainty 1.0), this change can be used to increase the relevance of this difference; this acts like an attention mechanism, gradually increasing the importance of this part of an agent's b-frame. Once this relevance passes a threshold, which can be indicated by another belief (perhaps unique to this b-frame), belief revision would begin. In short, we have a striking example of how one can utilize the integration of influence, certainty and relevance; a difference in *certainties* can be used to *influence* the *relevance* of this difference, which in turn can be used to trigger revision of any one of these three properties.

A brief example can illustrate the above strategy. A b-frame containing the beliefs $\text{certainty}(\text{agent}, B) = 0.5$ and $\text{certainty}(\text{agent}, \text{not}(B)) = 0.5$ would have a relevance too low to trigger belief revision. Intuitively, the above situation indicates a lack of an opinion about either belief. A different b-frame might have certainties of 0.8 and 0.2, respectively; these new numbers, like the previous ones, sum to 1.0 and hence might be considered non-contradictory, even though there isn't total support for either one. Note that this second b-frame might result after new evidence arrived and led to revision of the two beliefs' certainties. A third b-frame might have certainties of 0.8 and 0.8; this is getting too close to the standard contradiction scenario (i.e., 1.0 and 1.0), and hence might trigger belief revision.

These three b-frames hint that one informal rule of thumb might be to begin increasing the relevance of a belief pair X and not(X) if their total certainties deviate from 1. However, this rule might be fine for one agent (i.e., one b-frame) but replaced by a related yet different rule for another agent (i.e., a different b-frame). This is one illustration of the nature of belief relativity: that b-frames are best viewed as relative states, useful for adopting points of view and exploring the consequences of such views.

A General View of Belief Relativity

A wider view of BR leads to several interesting issues; specifically, one can produce certain claims based on an analogy between BR and the relativity theory (Einstein 1905) of the physical world (call it PR).

First, just as there is no absolute physical frame of reference from which measurements can be made with complete certainty, in BR *there is no absolute b-frame* from which to perform belief computations with complete certainty. For instance, just as velocity is a relative concept in PR, dependent on the frame one is observing from, the BR view is that *degree of belief is relative*, depending on what b-frame one is believing from. In PR, there exists the relativity of motion; driver A may claim "I'm travelling forward at 5 mph; car B is still", whereas driver B may claim "I'm travelling backward at 5 mph; car A is still". Both are correct, because constructing an answer is relative, and depends on the frame of reference from which one is viewing the situation.

An analogous example in BR is the case where one assertion can cause different degrees of belief to be computed in different b-frames – e.g., belonging to two different agents. For example, suppose an optimist O and a pessimist P view the same belief X representing the outcome of a mutual friend M's test: $X = (\text{grade}(M, \text{test}) = 79)$. O

may be happy about this: $\text{happiness}(O, X) = 0.8$, perhaps because this shows progress by M. P may not be as happy about the outcome: $\text{happiness}(P, X) = 0.5$, perhaps because P thinks M's parents will be upset that M didn't ace the test. Who is right? Both, because the answer is relative; multiple b-frames, even if they are constructing beliefs about the same property, can process the same belief with equal validity. That is, different b-frames can have different views of the same situation without contradicting each other. In this case, both happiness levels are "correct", depending on what b-frame one is in.

Note that the qualitative reasons cited above for O's and P's happiness levels might be based on other beliefs in their respective b-frames (i.e., previous experience concerning M's life). In general, BR provides a model for how one processes a belief a certain way (e.g., like an optimist or pessimist) based on the influence of other beliefs in one's b-frame. The same agent can even act like either an optimist or a pessimist, depending on what beliefs are used to compute degrees of certainty for a new belief (i.e., what b-frame one chooses to view a situation from). Finally, whereas this example dealt with how different b-frames view the same property (happiness), earlier examples showed how different b-frames applied to a belief can also utilize different property rules (e.g., tallness and relevance).

The second major claim is that *adopting a higher-level frame can resolve conflicts and increase certainty, but it can also have the opposite effect*. In PR, deciding which of two motions are relative can only be done by comparing them to objects in a larger frame. In the previous PR example, an observer using a higher (more omniscient) frame may decide that car A really was moving forward relative to car B, and that the reverse does not hold, because A is moving forward relative to other objects in the higher frame while B is standing still relative to them. Similarly, the previous BR example might change by adopting a higher frame which includes M's statement that he is as happy about his grade as O claimed; this resolves the relative views of O and P in favor of O. However, adopting still wider frames can again lead to belief revision. A higher frame in the car example might find that the object a believer used as a reference to compare to A and B (e.g., a train) is *itself* moving, and doing so with the same direction and speed as B; in the view of this higher b-frame, where the new reference object is the ground, A is now still while B and the train are moving, thus contradicting the earlier b-frame belief. Similarly, a wider b-frame in the BR example may find that M is lying about his happiness to hide his parents' unhappiness; this now resolves the O-P conflict in favor of P.

In summary, BR theory proposes that one can achieve *local certainty*, but this is the best one can do. Absolute certainty is unattainable, because there is always the potential of finding another frame that invalidates one's "omniscient" point of view. In short, one cannot attain permanent global absolute certainty, since there is no b-frame provable to be absolute over all beliefs for all time. However, finding temporary local certainty among one's beliefs is a useful goal, such as to resolve apparent contradictions by going to higher b-frame. In general, an optimal agent must be able to find b-frames that provide relative certainty, while also being smart enough to know when to search for more useful b-frames (either to supplant existing ones, or add new b-frames to one's repertoire). BR can hopefully provide a framework to model and improve such behaviors.

Related Work

In BR, generating new assertions is performed inside a b-frame, and evidence for these beliefs is also maintained there. In this sense, BR is directly related to systems such as TMS, ATMS, and REVOLVER. If one uses BR only to infer new assertions from old via rules, and all degrees of certainty are either zero or one, and relevance is not represented, then BR is essentially reduced to these frameworks. Like ATMS-based systems, BR systems settle for local coherence or certainty over global, and can opt to keep non-currently-active belief groups in memory. If one augments ATMS-based systems with a method for propagating fractional certainty levels, this type of system resembles influence diagrams, or belief networks. Further, if one uses hillclimbing search for consistent belief sets, this type of system would now resemble systems such as REVOLVER.

However, BR is designed to improve upon the above methods when reasoning about how different agents process beliefs, and how one agent can adopt different points of view. That is, BR provides a better mechanism for reasoning about multiple b-frames. For example, BR provides a stronger framework for dealing with nested beliefs – beliefs about the beliefs of others. Even other b-frames can be the focus of such nested belief (e.g., the degree of certainty agent X has in agent Y's beliefs about some topic). BR also can utilize a more complete set of revision operators than systems such as REVOLVER; the latter could search for and generate new premise beliefs, but BR also deals with creating new b-frames, as well as new influences among them.

Finally, BR utilizes an explicit notion of belief relevance that the systems mentioned above do not. This added concept enables relevance and certainty levels to influence each other to produce interesting behaviors (e.g., by nesting these concepts within one belief). However, degree of relevance can also be used to compute a degree of activeness; that is, relevance can be used to focus system attention on the most interesting assertions. An agent can use such knowledge to control the propagation of certainty degrees (e.g., perform such propagation only on those most-interesting assertions). In short, just as other areas of AI utilize knowledge to control search, so too can one use relevance information to control the search for useful beliefs (i.e., those that are relevant to current goals), the beliefs they influence, and the certainty levels of all such *belief families*. The above point is even more important due to the oft-cited problems TMS and ATMS systems have regarding excessive computation required to process even small sets of premises, due to storage of all inferred beliefs in memory. TMS and ATMS systems give beliefs a status of either active or inactive; TMS and ATMS can thus be viewed as a special case of BR where all relevance levels are either zero or one. By substituting degree of relevance for active/deactive status, BR beliefs can be prioritized for processing, which should lead to more efficient systems. As is the case for certainty levels, relevance levels can be calculated using a method (such as that of Thagard and Kunda (1987)) that will be independent of the methods used to maintain the influence networks.

The second useful comparison is to fuzzy systems (Zadeh 1965). We have seen how agents in BR can propagate degrees of certainty and relevance by using fuzzy membership functions. In addition, like fuzzy logic, BR can deal with the notion that an entity can "have" property P as well as

not(P). For example, one agent X can believe that another agent M is both happy and not happy in BR (where X is viewing M from a b-frame that encompasses those of two other agents observing M). However, unlike fuzzy logic, BR can deal with reasoning on the more abstract b-frame level, such as ATMS-type operations (e.g., maintaining chains of inference support), and modelling different points of view as well as what might trigger such changes.

Conclusion

I have presented a model called belief relativity, a model for reasoning about the structure of belief systems, and how degrees of certainty and relevance are inferred. BR provides a unified framework for ideas from previous research, such as influence/belief networks, TMS and ATMS systems, and hillclimbing systems for belief creation and revision.

An important advantage of BR is that influences, as well as the rules for propagating certainty and relevance values among them, can be represented and processed *independent* of each other. That is, BR belief systems can perform b-frame level tasks (keeping track of influences, trying out new sets of beliefs, creating new influences) yet also perform tasks complementary to those concerning b-frames (e.g., using membership functions to determine degrees of belief, and propagating them along the influences set up by b-frame reasoning). In short, BR should allow the best features from ATMS-style systems and fuzzy/probabilistic systems to be utilized in a synergistic manner.

References

- de Kleer, J. 1984. Choices without backtracking. *Proceedings of the Fourth National Conference on Artificial Intelligence* (pp. 79–85). Austin, TX: Morgan Kaufmann.
- Doyle, J. 1979. A truth maintenance system. *Artificial Intelligence*, 12, 231–272.
- Einstein, A. 1905. On the electrodynamics of moving bodies. *Annalen der Physik*, 17, 891.
- Festinger, L. 1957. *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.
- Howard, R. A. & Matheson, J. E. 1984. Influence Diagrams. In R. A. Howard and J. E. Matheson eds., *Readings on the Principles and Applications of Decision Analysis*, Vol. II, Strategic Decisions Group.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.
- Rose, D. 1989. Belief Revision and Machine Discovery. Ph.D. diss., Dept. of Information and Computer Science, University of California, Irvine.
- Rose, D., & Langley, P. 1986. Chemical discovery as belief revision. *Machine Learning*, 1, 423–451.
- Shafer, G. 1976. *A mathematical theory of evidence*. Princeton: Princeton University Press.
- Thagard, P. & Kunda, Z. 1987. Hot cognition: Mechanisms for motivated inference. *Proceedings of the Ninth Annual Conference of the Cognitive Science Society* (pp. 753–763). Seattle, WA: Lawrence Erlbaum Associates.
- Zadeh, L. A. 1965. Fuzzy sets. *Information and Control*, 8, 338–353.