

Qualitative Reasoning about Perception and Belief

Alvaro del Val*

Departamento de Ingeniería Informática Computer Science Department

Universidad Autónoma de Madrid

28049 Madrid, Spain

delval@ai.ii.uam.es

Pedrito Maynard-Reid II†

Stanford University

Stanford, CA 94305

pedmayn@cs.stanford.edu

Yoav Shoham

Computer Science Department

Stanford University

Stanford, CA 94305

shoham@cs.stanford.edu

Abstract

We present a qualitative model for reasoning about perceptions, sensors, and belief, and a logic to reason about this model. Basic to our model is a distinction between precision and accuracy, for both of which we provide qualitative definitions. In our logic this distinction gives rise to two modal operators— P for actual perception, and C_p for perceptual capability, which is captured as a set of possible percepts. Adding to these operators the standard B operator to model belief, we end up with a logic combining standard Kripke-style semantics with the almost-standard ‘neighborhood semantics.’ We define various agent types in the logic, from agents who believe all and only what the sensors tell them, to much more skeptical agents. We define each agent both axiomatically and model-theoretically, and provide soundness and completeness results relating the two types of definitions.

1 Introduction

A great deal of attention in AI has been devoted to modeling states of information, whether qualitatively through, e.g., logics of knowledge and belief (cf. [Moore, 1985; Konolige, 1986; Fagin *et al.*, 1995]), or quantitatively. In contrast, relatively little attention has been paid to modeling information *sources*, on the basis of which the various information states are reached. In particular, there has been very little work on the relationship between belief and *perception*, the most common source of information in everyday life. This is the topic of this paper. (We hasten to add that there certainly has been some work on this topic in AI, notably

by Davis [1989] and Bacchus, *et al.* [1995]; we discuss this work in section 5.)

The connection between belief and perception is multifaceted. Here we concentrate on a particular set of issues, namely, the precision and accuracy of percepts and of sensors that generate them. This distinction is important but easily glossed over; we discuss it in detail in the next section.

The contributions of this paper are as follows:

- A qualitative model of precision and accuracy.
- A logic with which to reason about certain aspects of the model. The logic contains three different but related modalities: P (for actual perception), C_p (for the set of possible percepts), and B (for belief). They are each defined in what by now is a standard fashion, but the combination is novel.
- A category of agent types, each defined both axiomatically and model-theoretically; the types differ either on the relationship between the agent’s percepts and the world, or (and this is the most interesting case) on the relationship between the agent’s percepts and his beliefs.
- Completeness results, relating the axiomatic systems and the semantic models.

The paper is organized as follows. In section 2 we present our model of precision and accuracy, and explain it with examples. In section 3 we present the logic **CP0** for modeling and reasoning about perception, and present our first completeness result. We also define the first two types of agent, *accurate* and *observant*. In section 4 we present the logic **BCP0** for modeling and reasoning about perception and belief, and present many more agent types. We provide completeness results for **BCP0** as well as for the various agent types. We discuss related work in section 5, and make some concluding remarks in section 6.

*This work performed at Stanford University’s Computer Science Department.

†This work partly supported by a National Physical Science Consortium Fellowship.

2 Qualitative model of accuracy and precision

Our starting point is a distinction between *precision* and *accuracy*, both between the precision and accuracy of a particular percept, and the precision and accuracy of a sensor.

We illustrate these notions through an example. Consider a sensor that is supposed to detect the location of an object on the real line. A percept (or reading) of this sensor would be some interval such as $[1, 7]$, meaning that the object is somewhere in that interval. Let's denote this percept by P_1 . Another percept might be $P_2 = [3, 5]$. In this case P_2 is said to be more precise than P_1 . Suppose that the *actual* location is 2; then P_1 is said to be more accurate than P_2 since P_1 is right and P_2 is wrong. Now consider $P_3 = [3, 4]$. Clearly P_3 is more precise than P_2 , but which of them is more accurate? The intuitive answer is not obvious, since they are both wrong, but in our model P_3 will come out more accurate. Intuitively, P_2 allows locations that are "more wrong" than does P_3 . Note that in order to capture "more wrong" we need to appeal to the notion of a "similarity ordering"—which of two points are closer to a third—but not necessarily an exact metric specifying the distance between two points.

The three example percepts discussed so far are related by set inclusion, but now consider $P_4 = [4, 6]$; is it more or less precise than P_2 ? The answer depends on whether we can appeal to metric properties or not. If we can then we can judge the two percepts equally precise. However, since we wish to capture qualitative notions, in our model the precisions of P_2 and P_4 will come out as non-comparable. But what about their accuracies? In our model P_2 will come out as more accurate; the intuition is again that P_4 allows locations that are even more wrong than does P_2 .

Let us now begin to develop a formal model; in the following definitions assume a set S (of worlds).

Definition 1 A percept is a subset of S .

Definition 2 Percept X is said to be as precise as percept Y iff $X \subseteq Y$.

Definition 3 Percept X is accurate at $s \in S$, abbreviated *s-accurate*, iff $s \in X$.

Definition 4 Given a set of allowable percepts A , percept X is as accurate as percept Y at $s \in S$ wrt A iff $\forall Z \in A$, if $s \in Z$ and $Y \subseteq Z$ then $X \subseteq Z$.

Thus X is as accurate as Y iff every accurate percept which includes Y must also include X . This definition is intuitive but often quite weak. Therefore we consider ways to impose more structure upon possible perceptions, in particular a notion of "convexity" of the space derived from a similarity preorder.

Definition 5 A similarity ordering on S at $s \in S$ is a binary relation $\leq_s \subseteq S \times S$; the intuitive reading of $\leq_s(s_1, s_2)$, which we abbreviate $s_1 \leq_s s_2$, is " s_1 is at least as close as s_2 to s ". $<_s$ denotes the strict version of $\leq_s$. We require that $\leq_s$ be a preorder (that is, reflexive and transitive), and that if $s \neq s'$ then $s <_s s'$.

Definition 6 Given $\leq_s$ as above, percept X is convex at $s \in S$, abbreviated *s-convex*, iff $\forall x, y \in X$, if $x \leq_s z \leq_s y$ then $z \in X$. A set of percepts A is *s-convex* iff every $X \in A$ is *s-convex*.

Definition 7 The accurate extension of a percept X at $s \in S$ is the percept $X^+(s) = \{y \in S \mid \exists x \in X : y \leq_s x\}$.

Definition 8 A set of percepts A is strongly *s-convex* iff A is convex and $\forall X \in A : X^+(s) \in A$.

Proposition 1 $X^+(s)$ is *s-convex* and *s-accurate*, and for any *s-convex* and *s-accurate* percept $Z \in A$, if $X \subseteq Z$ then $X^+ \subseteq Z$.

In the following propositions, assume A to be strongly *s-convex* and $X, Y \in A$.

Proposition 2 X is as accurate as Y at $s \in S$ wrt A iff $X^+(s) \subseteq Y^+(s)$.

Proposition 3 X is as accurate as Y at s wrt A iff $\forall x \in X \exists y \in Y : x \leq_s y$.

Proposition 4 Let $\leq_s$ be a total preorder. X is as accurate as Y at $s \in S$ wrt A iff $\sup(X) \leq_s \sup(Y)$, where $\sup(Z)$ is an element of set Z that is a least upper bound wrt $\leq_s$.

Under the assumptions of this last proposition, therefore, the relation "being as accurate as" is a total preorder. Also, if A is convex but not strongly convex, then the right hand sides of these three propositions are sufficient but not necessary for X to be as accurate as Y .

So far we have discussed the relative precision and accuracy of percepts; let us now turn to comparing sensors. We model a sensor by its *perceptual capabilities*, which, loosely speaking, means the set of percepts it is capable of delivering in any particular circumstance. Because sensors represent sets of percepts, their relative precision and accuracy can be defined in a variety of ways depending on the application. Two possibilities are as follows: (1) Sensor A is considered more precise/accurate than sensor B in some situation if every percept of which sensor A is capable is more precise/accurate than some percept sensor B could deliver. (2) Sensor A is more precise/accurate than sensor B in a situation if every percept sensor A could return is more precise/accurate than any percept sensor B could return. The first case corresponds to a comparison of the worst case scenarios of the two sensors, whereas the second case corresponds to a guarantee on the relative precision/accuracy of the percepts actually returned. The second is obviously more restrictive, and implies the first.




Figure 1: Two-sensor example.

Again, we illustrate these notions through an example. We will also demonstrate that not only can a given sensor's capabilities change with the world, but so can the relative relationship between the capabilities of different sensors. Consider two sensors, A and B, which report on the location of some object on a line. Imagine the two sensors as located on another line (the *baseline*) parallel to the line the object is on (see Figure 1).

The capabilities of sensor x (where x is A or B) are characterized by two parameters, α_x and γ_x . α_x is the angle of resolution of sensor x , whereas γ_x is the maximum angle of divergence between the orientation implied by the particular reading and the sensor's actual orientation. Given these parameters, the distance of the baseline from the real line, and the location and orientation of sensor x , we can compute δ_x , the size of the interval it returns, and ϵ_x , the maximum displacement of the reading from the true interval. (We assume that γ_x is small enough compared to α_x that the variation in interval length due to γ_x is negligible.) Let us suppose that $\alpha_B < \alpha_A$ and $\gamma_B < \gamma_A$.

The object's location determines the world, and at a world each sensor returns one percept. To determine the percept it returns at a world, sensor A maintains a fixed orientation perpendicular to the real line, and travels back and forth along the baseline, periodically taking a reading until the object appears in its vision. Because its orientation is constant (within γ_A), both δ_A and ϵ_A are constant. Imagine, however, that it is more expensive for sensor B to move around, so instead it remains in one location and varies its orientation by rotating, randomly sampling the real line until the object appears in its vision. In this case both δ_B and ϵ_B vary with the sensor's orientation, so both its accuracy and precision decrease the farther away the object is.

Thus the specific set of possible percepts which sensor x might deliver when the object is located at l consists of the intervals $[r, r + \delta_x]$ for all $l - \delta_x - \epsilon_x \leq r \leq l + \epsilon_x$

(corresponding to the extreme cases where the the interval is ϵ_x to the left and right of the object, respectively) for $r \in \mathbb{R}$ (the real numbers). Consider the world where the object is located directly in front of sensor B. Sensor B is more precise than sensor A because all the percepts sensor B might return are more precise than all those sensor A might return. It is also more accurate (under the worst-case interpretation) than sensor A because the least accurate percept it is capable of is more accurate than some of sensor A's possible percepts.

Now suppose the object is far away from sensor B. The size of the interval returned by sensor B and the magnitude of its maximum displacement will be much larger. In this case sensor B will be both less precise and less accurate (in the worst case) than sensor A because it will have possible percepts that are less precise and less accurate than all the percepts sensor A might return.

The definitions below extend our framework to include sensors. Assume a set of allowable percepts A as before.

Definition 9 A sensor is a function $f : S \rightarrow 2^{2^S}$ that describes the set of possible percepts in a world.

Definition 10 Sensor f is worst-case as precise as sensor f' at $s \in S$ iff $\forall X \in f(s) \exists Y \in f'(s)$ s.t. X is as precise as Y at s . It is guaranteed as precise as sensor f' at s wrt A iff $\forall X \in f(s) \forall Y \in f'(s)$, X is as precise as Y at s wrt A .

Definition 11 Sensor f is worst-case as accurate as sensor f' at $s \in S$ wrt A iff $\forall X \in f(s) \exists Y \in f'(s)$ s.t. X is as accurate as Y at s wrt A . It is guaranteed as accurate as sensor f' at s wrt A iff $\forall X \in f(s) \forall Y \in f'(s)$, X is as accurate as Y at s wrt A .

3 A logic for perception

We now aim to construct a logic with which to reason about these aspects of perception. As the above discussion illustrates, at a basic level reasoning about perception involves the counterfactual notion of what *could be perceived*, and it is this notion we will concentrate on. We capture it formally by introducing a logic with two modal operators, P and C_p , which represent actual perception and possible perception, respectively.

This gives us a way to capture the perceptual limitations of an agent. Next we want to represent the degree to which the agent is aware of these limitations. What an agent believes given some percept may depend on the agent's level of trust in the reliability of her sensors, which in turn may depend to some degree on the actual percepts themselves. To be able to describe complex interactions of this sort between the agent's mental state and her perceptions and possible perceptions, we also add to the formalism a modal operator, B , for belief. We then use this model to define several agent types with various interesting perceptual-belief characteristics.

3.1 The logic CP0

The formulas of the language $\mathcal{L}$ are obtained by closing a set of propositional symbols under the boolean connectives, and under the operators P and C_p . Intuitively $P\phi$ means “the agent perceives ϕ ”, where ϕ is some arbitrary formula. $C_p\phi$ indicates that “the agent can perceive ϕ ”, meaning that her perceptual apparatus could have produced a percept that validates ϕ . We also ensure, of course, that what is perceived can be perceived.

For semantics, we define a P-structure as a tuple $\mathcal{M} = (S, N, I, \pi)$ where S is a set of states or worlds, π is a truth assignment $\pi : S \rightarrow 2^{\mathcal{L}}$, N is a relation $N : S \rightarrow 2^{2^S}$ mapping each world to a set of sets of worlds, and I is an accessibility relation $I : S \rightarrow 2^S$ mapping each world to a set of worlds.

The satisfaction relation is defined as usual. Only the truth conditions for the operators P and C_p are of interest. They are:

- $\mathcal{M}, s \models C_p\phi$ iff $\exists Z \in N(s)$ s.t. $\forall t \in Z, \mathcal{M}, t \models \phi$,
- $\mathcal{M}, s \models P\phi$ iff $\forall t \in W$, if $t \in I(s)$ then $\mathcal{M}, t \models \phi$.

Thus P uses the standard semantics for necessity operators in modal logic, whereas C_p uses a variant of neighborhood (or minimal model) semantics [Chellas, 1980]. For $s \in S$, we use the notation $N(s) = \{X \subseteq S \mid (s, X) \in N\}$, and $I(s) = \{t \in S \mid (s, t) \in I\}$. The idea is that $N(s)$ represents the set of possible percepts at a world, and the set A of allowable percepts referred to in the previous section includes all these possible percepts from all worlds. In a domain that is strongly s -convex, A would also include all the accurate extensions of these percepts, although these extensions may not be actually perceivable from any world.

We are interested in the class $\mathcal{M}_{CP}$ of P-structures for perception satisfying the following conditions:

- CP1. $N(s) \neq \emptyset$.
- CP2. $\emptyset \notin N(s)$.
- CP3. $I(s) \in N(s)$

Condition (CP3) is the most interesting one. It requires every actual percept to be a possible percept, and corresponds to axiom A4 below. (CP1) and (CP2) correspond, in turn, to axioms A2 and A3, respectively.

Let **CP0** be the following set of axioms and rules of inference:

- A1. All tautologies of propositional calculus.
- A2. $C_p \text{true}$
- A3. $\neg C_p \text{false}$
- A4. $P\phi \supset C_p\phi$
- A5. $(P\psi \wedge P(\psi \supset \phi)) \supset P\phi$.
- R1. From ψ and $\psi \supset \phi$ infer ϕ .

R2. From $\psi \supset \phi$ infer $C_p\psi \supset C_p\phi$.

R3. From ψ infer $P\psi$.

Theorem 1 CP0 is a sound and complete axiomatization of the class $\mathcal{M}_{CP}$ of P-structures for perception.

Proposition 5 These formulas are valid in $\mathcal{M}_{CP}$:

- 1. $C_p(\phi \wedge \psi) \supset (C_p\phi \wedge C_p\psi)$.
- 2. $\neg P \text{false}$.

Thus **CP0** includes only some weak consistency conditions on C_p , namely A2 and A3, and on P , in addition to the aforementioned restriction that actual percepts must be possible percepts. On the other hand, the system does *not* sanction the axiom $(C_p\psi \wedge C_p\phi) \supset C_p(\psi \wedge \phi)$: for example, it may be possible for me to perceive some object as either blue or red depending on lightning conditions, but this does not mean that I can perceive the object as both red and blue.

3.2 Accurate and observant agents

Various agent types can be defined to capture varying assumptions about the application domain. Most definitions refer to the relationship between perception and belief, a topic we have not yet addressed. However, some types are defined in terms of the relationship between perception and objective facts in the world, independent of the belief modality. Thus we give them here.

We first consider *observant* agents, those from whose perception nothing escapes. Formally, the agent is said to be observant if the following schema is valid:

$$\phi \supset P\phi.$$

Similarly, the agent is said to be *accurate* if the schema

$$P\phi \supset \phi$$

is valid.

The schema corresponding to the observant agent is clearly very strong, unlikely to be applicable to agents in more complex environments, since nothing perceptually escapes the agent. Yet it is possible to think of simple agents for which it makes sense, such as, for example, a device continuously and accurately monitoring some variable, and whose model of the world consists exclusively of the possible values the variable can take.

We can also consider similar properties for the C_p modality. The agent is *potentially observant* iff she can perceive any (true) fact, i.e. iff the following schema is valid:

$$\phi \supset C_p\phi.$$

An agent is said to be *necessarily accurate*, on the other hand, if anything she can perceive is true, i.e. iff the following schema is valid:

$$C_p\phi \supset \phi.$$

4 A logic for perception and belief

We are now interested in what the agent should believe, given what she perceives. It might be tempting to make the assumption that anything perceived is also believed. However, in (common) cases where the sensors might not be completely accurate, and in which the agent might know this, the assumption would be inappropriate. Sometimes the converse assumption is more adequate: one doesn't believe in anything without having perceived it. Keeping this in mind, we begin with the most general case, in which the relation between perception and belief is unconstrained, and which can be captured in a straightforward manner. We will then identify a number of special cases with further properties.

4.1 The logic BCP0

We expand CP0 to include an operator B . A PB-structure is a tuple $\mathcal{M} = (S, N, I, R, \pi)$, where S , N , and I and π are as before (in particular (CP1)–(CP3) are satisfied), and R is a serial, transitive and Euclidean relation over S . That is, we are characterizing belief with the standard KD45 system. In addition, we require the agent to be aware of her perceptions, by imposing the following semantic condition (which corresponds to axioms A10 and A11 below):

BCP1. $\forall s, t$, if $t \in R(s)$ then $I(s) = I(t)$.

Let $\mathcal{M}_{BCP}$ be the class of such structures. The semantics of the B operator are given as usual by:

$\mathcal{M}, s \models B\psi$ iff $\forall t \in S$, if $t \in R(s)$ then $\mathcal{M}, t \models \psi$.

The system **BCP0** consists of all the axioms and rules of **CP0** together with:

A6. $(B\psi \wedge B(\psi \supset \phi)) \supset B\phi$.

A7. $\neg B(\text{false})$.

A8. $B\psi \supset BB\psi$.

A9. $\neg B\psi \supset B\neg B\psi$.

A10. $P\psi \supset BP\psi$.

A11. $\neg P\psi \supset B\neg P\psi$.

R4. From ψ infer $B\psi$.

The following result should be obvious:

Proposition 6 *BCP0 is a sound and complete axiomatization of the class $\mathcal{M}_{BCP}$ of PB-structures for perception and belief.*

Note that the formula $BP\psi \equiv P\psi$ follows from A10 and A11, and is therefore a theorem of **BCP0**. Although one can easily think of counterexamples, e.g. ambient noise, or sensory overload, where percepts are filtered out so that the agent is not aware of what she is perceiving, we have elected to adopt this assumption to simplify the

remaining discussion. It is easy to weaken the formalism to deal with this situation if desired.

Given this logic, we can make additional inferences about perception and belief. For example, the formula

$$P\phi \wedge B(\psi \supset \neg C_p\phi) \supset B\neg\psi$$

is valid in **BCP0**. We can read this as saying that if the agent perceives ϕ and she believes she could not have perceived that in a world in which ψ holds, then she believes that ψ is false. We can obtain stronger results by imposing additional conditions on the agent. This motivates the definition of a number of agent types according to the connections between the various modalities. We do this in the next section.

4.2 More agent types

By placing additional joint restrictions on I , N , and R , it is possible to capture different ways in which perceptions and perceptual capabilities can be related to the agent's beliefs and to the environment. There are a number of basic dimensions to consider. In a previous section we appealed to the relationship between perception and the environment in defining the accurate and observant agents, and also between perceptual capabilities and the environment; we also decided to build into the logic the introspective abilities of the agent with respect to perception through axioms A10 and A11. In addition we may now consider the introspective abilities of the agent with respect to her perceptual capabilities, and the confidence that the agent places in her perceptions. As was discussed in the previous section, a third dimension—the correctness of the agent's beliefs about her perceptions—is built into the logic. Table 1 defines the various agent types along these dimensions in terms of the validity of certain (not necessarily independent) axiom schemas, and Table 2 identifies certain classes of PB-structures for perception and belief in terms of restrictions on I , N , and R . The correspondence between the axiom schemas and the corresponding classes of structures is given in the soundness and completeness results in the next theorem, where **BCP0**+ i refers to the system composed of **BCP0** and the i th axiom schema in Table 1.

Theorem 2 *BCP0 + 1 (respectively + 2, + 3¹, + 4, + 5, + 6, + 7, + 8, + 9) is a sound and complete axiomatization of the class of PB-structures for perception and belief $\mathcal{M}_{\text{observant}}$ (respectively $\mathcal{M}_{\text{accurate}}$, $\mathcal{M}_{\text{Cp-observant}}$, $\mathcal{M}_{\text{Cp-accurate}}$, $\mathcal{M}_{\text{Cp-aware}(+)}$, $\mathcal{M}_{\text{Cp-aware}(-)}$, $\mathcal{M}_{\text{confident}}$, $\mathcal{M}_{\text{cautious}}$, $\mathcal{M}_{\text{sib.}}$)*

¹The completeness proof for the potentially observant agent assumes a finite propositional vocabulary.

Agent Type	Axiom schema
Observant agent	1. $\psi \supset P\psi$
Accurate agent	2. $P\psi \supset \psi$
Potentially observant	3. $\psi \supset C_p\psi$
Necessarily accurate	4. $C_p\psi \supset \psi$
Aware (+) of perceptual capab.	5. $C_p\psi \supset BC_p\psi$
Aware (-) of perceptual capab.	6. $\neg C_p\psi \supset B\neg C_p\psi$
Perceptually confident	7. $P\psi \supset B\psi$
Cautiously confident	8. $P\psi \supset \neg B\neg\psi$
SIB (seeing-is-believing)	9. $\neg P\psi \supset \neg B\psi$

Table 1: Agent types: axiomatic conditions.

Class of structures	Restriction
$\mathcal{M}_{\text{observant}}$	1. $\forall s, t : t \in I(s) \Leftrightarrow s = t$
$\mathcal{M}_{\text{accurate}}$	2. $\forall s : s \in I(s)$
$\mathcal{M}_{C_p\text{-observant}}$	3. $\forall s : \{s\} \in N(s)$
$\mathcal{M}_{C_p\text{-accurate}}$	4. $\forall s : s \in \bigcap N(s)$
$\mathcal{M}_{C_p\text{-aware}(+)}$	5. $\forall s, t : t \in R(s) \Rightarrow N(s) \subseteq N(t)$
$\mathcal{M}_{C_p\text{-aware}(-)}$	6. $\forall s, t : t \in R(s) \Rightarrow N(t) \subseteq N(s)$
$\mathcal{M}_{\text{confident}}$	7. $\forall s : R(s) \subseteq I(s)$
$\mathcal{M}_{\text{cautious}}$	8. $\forall s : R(s) \cap I(s) \neq \emptyset$
$\mathcal{M}_{\text{sib}}$	9. $\forall s : I(s) \subseteq R(s)$

Table 2: Agent types: semantic conditions.

5 Related work

There has been relatively little work relating perception and belief; here are the closest of which we are aware.

First, related to our paper is Davis' earlier work [Davis, 1989] on belief and perceptual acuity. Davis' work, however, like some other work in related fields (cf., [Parikh, 1983]), is firmly based on modeling metric spaces, and is thus quantitative in nature. Davis' logic contains only notions analogous to our P and B , and does not have a way of specifying the perceptual capabilities of the sensor. He was thus not able to differentiate in the logic between precision and accuracy, and ran into some paradoxes as a result.

Recently, Bacchus, *et al.* [1995] have worked on representing the effects of readings from noisy sensors on the belief state and subsequent actions of agents. They employ a formalism that combines a probabilistic definition of belief with a situation calculus framework. However, our approach, while remaining amenable to probabilistic extensions, is able to capture qualitative effects of perception as well. Their work also does not address precision.

6 Discussion

We have presented a formalism to reason about perception and belief, an under-explored area of investigation. We started by distinguishing between precision and ac-

curacy, notions that are often conflated. We proposed what to our knowledge is the first qualitative definition of these notions.

We then provided a model in which to represent an agent's perceptions and beliefs, and a logic with which to reason about the model. We explored possible relations between an agent's perceptions, perceptual capabilities, and beliefs, and between these and the environment. Because no type of agent among those defined is likely to be useful in all scenarios, we have mentioned some circumstances in which different types might be used.

The formalism we have presented seems natural and well-behaved, but more research is needed on applying it, whether in general formulations of common-sense reasoning or in more specific applications. It would also be interesting to study extensions of the language that incorporate more specific domain structure (such as a metric space), and relate them to perceptual indistinguishability. Temporal and multi-agent aspects are also of interest, as well as the problem of updating the beliefs of the agent in the light of perceptually acquired information. Indeed, we view sensors as a source of information, and precision/accuracy as measures of that source's competence. We intend to apply this theory more generally to reason about multiple, partially-trusted information sources.

Acknowledgements We would like to thank the members of the Nobots research group at Stanford University for their helpful comments, as well as the referees for constructive critique on the paper.

References

- [Bacchus *et al.*, 1995] F. Bacchus, J. Halpern, and H. Levesque. Reasoning about noisy sensors in the situation calculus. In *Proc. 14th Conf. IJCAI*, pages 1933–1940, 1995.
- [Chellas, 1980] B. Chellas. *Modal Logic: An Introduction*. Cambridge University Press, 1980.
- [Davis, 1989] E. Davis. Solutions to a paradox of perception with limited acuity. In *Proc. of the 1st Intl. Conf. on Knowl. Rep. and Reas.*, 1989.
- [Fagin *et al.*, 1995] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
- [Konolige, 1986] K. Konolige. *A Deduction Model of Belief*. Pittman Research Notes in AI, 1986.
- [Moore, 1985] R. Moore. Semantical considerations on non-monotonic logic. *Art. Int.*, 25:75–94, 1985.
- [Parikh, 1983] R. Parikh. The problem of vague predicates. In R. Cohen and M. Wartofsky, editors, *Language, Logic, and Method*, pages 241–261. D. Reidel, 1983.