

Focused Belief Revision as a Model of Fallible Relevance-Sensitive Perception

Haythem O. Ismail and Nasr Kasrin

German University in Cairo
Department of Computer Science
{haythem.ismail,nasr.kasrin}@guc.edu.eg

Abstract. We present a framework for incorporating perception-induced beliefs into the knowledge base of a rational agent. Normally, the agent accepts the propositional content of perception and other propositions that follow from it. Given the fallibility of perception, this may result in contradictory beliefs. Hence, we model high-level perception as belief revision. Adopting a classical AGM-style belief revision operator is problematic, since it implies that, as a result of perception, the agent will come to believe everything that follows from its new set of beliefs. We overcome this difficulty in two ways. First, we adopt a belief revision operator based on relevance logic, thus limiting the derived beliefs to those that relevantly follow from the new percept. Second, we focus belief revision on only a subset of the agent's set of beliefs—those that we take to be within the agent's current focus of attention.

1 Introduction

Agent X is about to cross the street. It turns left, and it sees a car approaching fast. X immediately stops and decides to wait until the car had passed and then retry. X 's decision to stop is clearly based on reasoning with background practical knowledge. If not for this knowledge, and this reasoning, X would have been in trouble. Now consider agent Y .

Y is about to cross the street. It turns left, and it sees a car approaching fast. Trusting in the veridicality of its percept, Y comes to believe that a car is approaching fast. Unlike X , however, Y never stops. Y is certainly in trouble, but is it to blame? If Y does not have the same practical knowledge that saved X , then whatever it did, though a sign of ignorance, is rational. However, if Y did have the valuable knowledge, but did not use it to derive useful conclusions, then it certainly is to blame. Finally, consider agent Z .

Z is about to cross the street. It turns left, and it sees a car approaching fast. Z immediately stops and decides to wait until the car had passed and then retry. Z stops walking, but it does not stop reasoning about the approaching car. Assuming it has common knowledge of automobiles, Z will infer that *this* car has a steering wheel, an accelerator, and a battery, and that, since it is moving, it must be under the control of a driver. Moreover, it will also infer that the driver is a person, and that, like all people, the driver probably has two hands

and two feet, with five digits in each. Z never stops reasoning, and never crosses the street.

Evidently, perception involves some element of reflection on what is perceived. Neither no reflection nor unbounded reflection are appropriate, as the cases of Y and Z attest. Let us refer to this kind of perception-induced reasoning, or reflection, as “high-level perception”. In this paper, we present a framework for high-level perception within a grounded, layered agent architecture [1,2]. In this architecture, perception results in the agent’s acquiring a new belief of the form “I now perceive s through m ”, where s is a perceivable state and m is one of the agent’s perceptual modalities. As defined above, high-level perception is not the mere addition of such a belief to the agent’s belief store; normally, the agent will also come to believe that s and other states (that follow from it) hold. But this might result in the agent’s holding contradictory beliefs. Hence, we model high-level perception as belief revision, which is compatible with belief update in our temporal language. Adopting a classical AGM-style belief revision operator satisfying deductive closure is problematic [3], since it implies that, as a result of perception, the agent will come to believe everything that follows from its new set of beliefs. We overcome this difficulty in two ways. First, we adopt a belief revision operator based on relevance logic [4], thus limiting the derived beliefs to those that relevantly follow from the new percept. Second, we focus belief revision on only a subset of the agent’s set of beliefs—those that we take to be within the agent’s current focus of attention. This allows us to avoid Z -like agents.

Related work on knowledge representation aspects of perception is rather slim. Most of this work presents multi-modal logics of the interactions between perception and belief [5,6,7]. Wooldridge and Lumosio describe a multi-modal logic \mathcal{VSK} of visibility, perception, and knowledge [8]. The semantics is based on automata-like models of agents and environments rather than possible worlds. All these systems, however, have nothing to say about the issue of high-level perception as we described it above. They also have little to say about the link between perception and belief revision. Our notion of focused belief revision is related, but not identical, to the *local revision* of [9]. Limitations of space prevent us from a careful comparison of the two notions. Suffice it to say that our focused revision seems to subsume local revision and to be more suitable to the treatment of perception.

2 From Sensing to Believing

The theory of agents adopted here is based on the GLAIR agent architecture [1,2]. GLAIR is a layered architecture consisting of three levels. The bottom level, the sensory-actuator level (SAL), is the level controlling the operation of sensors and actuators (being either hardware or simulated). The top level, the knowledge level (KL), is the level at which conscious symbolic reasoning and planning take place. This level is implemented by the SNePS knowledge representation, reasoning, and acting system [10]. Statements in the SNePS knowledge base

are taken to represent the beliefs of the agent. The SNePS reasoning system is a natural deduction system of relevance logic [4,11]. The middle level, the perceptuo-motor level (PML), provides an interface between the SAL and the KL for both perception and action. This is the level at which recognition (and cognition) of feature-vector representations of external objects and states takes place. The result of this activity is the formation of a belief to be added to the top KL. (See [2] for details.)

Since the details of the language used to represent the agent's beliefs are orthogonal to the main message of this paper, we minimize the exposition of its syntax and semantics. We assume a first-order language \mathcal{L} , with a rich ontology including individuals, time points, acts, and states; states may be thought of as propositional fluents of the situation calculus. The exact structure of time has no bearing on the results we present, but we may assume an unbounded, discrete, linear structure. A sentence of the form $Holds(s, t)$ means that state s holds at time t . A functional term of the form $Prog(a)$ denotes the state that holds whenever act a is in progress. For every perceptual modality m of the agent, we shall have a predicate symbol P_m , where a sentence $P_m(s, t)$ states that the agent has a perceptual experience of state s (or, simply, perceives s) at time t .

In the rest of the paper, by "perception beliefs" we mean beliefs that emanate from the PML into the KL as a result of a perceptual experience. We represent perception beliefs by sentences of the form $P_m(s, *NOW)$. Such beliefs are formed at the PML, where s is a perceivable-state term whose sub-terms and function symbols are all grounded in PML representations; $*NOW$ denotes the current time. (See [2].) A perception theory is a logical theory in \mathcal{L} that allows us to infer from a perception belief $P_m(s, *NOW)$ the belief $Holds(s, *NOW)$. Such an inference is clearly defeasible, since perception need not be veridical. Both Musto and Konolige [6] and Bell and Huang [7] present non-monotonic modal logics of perception and belief to account for this defeasibility. Their central claim is that, unless the agent has reason to believe that the perception process is abnormal, it should come to believe in the content of its perception. Abnormalities in perception may be due to faulty sensors, illusory environments with abnormal perception conditions, or hallucination. A sophisticated-enough agent will have a theory of its own perceptual apparatus, of the different ways illusions may arise for each of its modalities, and of hallucination. These theories are necessarily causal as argued in [6,7]. (See [7] for an example of a simple theory of the effects of red light on visual perception, and [12] for a theory of noisy sensors.) Assuming such a theory, for each modality m , we will have a perception-belief axiom of the form

$$\mathbf{PBm.} \quad \forall s, t [P_m(s, t) \wedge \mathcal{Y}_m(s, t) \supset Holds(s, t)]$$

where $\mathcal{Y}_m(s, t)$ is a sentence stating that the perception of s at t via modality m is normal. The set of these axioms, together with theories of perceptual apparatus, illusions, and hallucinations form what we call a perception theory. We will, henceforth, refer to this theory as Π .

Now, if the agent does not believe $\mathcal{Y}_m(s, t)$, it will not succumb to believing $Holds(s, t)$, unless it has other reasons. On the other hand, if the agent believes

$\mathcal{Y}_m(s, t)$, it will accept $Holds(s, t)$. This, however, might result in a contradiction, if the agent has independent reasons for rejecting $Holds(s, t)$. In such a case, the agent should revise its beliefs. In neither [6] nor [7] do we find a complete discussion of how this is achieved. (In [6], we find no discussion.) In the following section, we present a belief revision operator, focused belief revision, that we use in Section 4 to model high-level perception.

3 Focused Belief Revision

We assume a proof theory based on Anderson and Belnap’s system FR of relevant implication [4]. Cn_R will be henceforth used to denote relevance logic consequence. In FR , the relevant implication operator \Rightarrow is defined in such a way that, where A is a set of formulas, $\phi \Rightarrow \psi \in Cn_R(A)$ if ϕ is actually instrumental to the derivation of ψ from A . We follow the presentation in [13] of an assumption-based reason maintenance system that implements belief revision in SNePS [11,14]. Incidentally, assumption-based reason maintenance requires the recording of hypotheses used in derivations, which is independently motivated by the use of relevance logic.

Definition 1. A *support set* of a sentence $\phi \in \mathcal{L}$ is a set $s \subseteq \mathcal{L}$ such that $\phi \in Cn_R(s)$. s is *minimal* if, for every $s' \subset s$, $\phi \notin Cn_R(s')$.

Definition 2. A *belief state* \mathbb{S} is a quadruple $\langle \mathcal{K}, \mathcal{B}, \sigma, \preceq \rangle$, where:

1. $\mathcal{K} \subseteq \mathcal{L}$ is a *belief set*.
2. $\mathcal{B} \subseteq \mathcal{K}$, with $\mathcal{K} \subseteq Cn_R(\mathcal{B})$, is a finite *belief base*. If $\phi \in \mathcal{B}$, then ϕ is a *base belief*.
3. $\sigma : \mathcal{L} \rightarrow 2^{2^{\mathcal{B}}}$ is a *support function*, where each $s \in \sigma(\phi)$ is a minimal support set of ϕ . Further, $\sigma(\phi) \neq \emptyset$ if and only if $\phi \in \mathcal{K}$. In particular, if $\phi \in \mathcal{B}$, then $\{\phi\} \in \sigma(\phi)$.
4. $\preceq \subseteq \mathcal{B} \times \mathcal{B}$ is a total pre-order on base beliefs.

Base beliefs are beliefs that have independent standing; they are not in the belief state based *solely* on inference. In particular, perception beliefs (in the sense of Section 2) are necessarily base beliefs. The belief set \mathcal{K} is not closed under Cn_R ; it represents the set of sentences that are either base beliefs or that were *actually derived* from base beliefs. This is in contrast to the logically-closed $Cn_R(\mathcal{K})$ which is the set of sentences *derivable* from base beliefs.

The set $\sigma(\phi)$ is the family of minimal support sets that were actually used, or discovered, to derive ϕ . \mathcal{B} may include minimal support sets of ϕ that are, nevertheless, not in $\sigma(\phi)$, if they are not yet discovered to derive ϕ . The total pre-order \preceq represents a preference ordering over base beliefs. We will refrain from making any commitments about the origins of this ordering; for the purpose of this paper, the ordering is just given. Note that, given the finiteness of \mathcal{B} , \preceq is guaranteed to have minimal and maximal elements.

For brevity, where $\phi \in \mathcal{L}$ and $A \subseteq \mathcal{L}$, let $Cn_R(A, \phi) = \{\psi \mid \phi \Rightarrow \psi \in Cn_R(A)\}$. In what follows, $\mathbb{S} = \langle \mathcal{K}, \mathcal{B}, \sigma, \preceq \rangle$ is a belief state, $\mathcal{F} \subseteq \mathcal{K}$, and $\phi \in \mathcal{L}$.

Definition 3. A *focused expansion* with *focus set* \mathcal{F} of \mathbb{S} with ϕ is a belief state $\mathbb{S} \dot{+}^{\mathcal{F}} \phi = \langle \mathcal{K}_{+\mathcal{F}\phi}, \mathcal{B}_{+\mathcal{F}\phi}, \sigma_{+\mathcal{F}\phi}, \preceq_{+\mathcal{F}\phi} \rangle$, satisfying the following properties.

- (A⁺1) *Success:* $\mathcal{B}_{+\mathcal{F}\phi} = \mathcal{B} \cup \{\phi\}$.
- (A⁺2) *Relevant inclusion:* $\mathcal{K}_{+\mathcal{F}\phi} = \mathcal{K} \cup \text{Cn}_R(\mathcal{F}, \phi)$.
- (A⁺3) *Relevant Support:* For every $\psi \in \mathcal{L}$,
 1. $\sigma(\psi) \subseteq \sigma_{+\mathcal{F}\phi}(\psi)$; and
 2. if $\sigma(\psi) \subset \sigma_{+\mathcal{F}\phi}(\psi)$ then $\psi \in \text{Cn}_R(\mathcal{F}, \phi)$ and for every $s \in \sigma_{+\mathcal{F}\phi}(\psi) \setminus \sigma(\psi)$, there is an s'' such that $s \in \{s'' \cup s' \mid s' \in \sigma_{+\mathcal{F}\phi}(\phi)\} \subseteq \sigma_{+\mathcal{F}\phi}(\psi)$.
- (A⁺4) *Order preservation:* $\preceq_{+\mathcal{F}\phi}$ is a total pre-order on $\mathcal{B}_{+\mathcal{F}\phi}$ such that, for every $\psi, \xi \in \mathcal{B}$, $\psi \preceq_{+\mathcal{F}\phi} \xi$ if and only if $\psi \preceq \xi$.

The belief state resulting from focused expansion by ϕ will include ϕ and anything that follows from it, given the focus set \mathcal{F} . That all newly derived sentences indeed follow from ϕ is guaranteed by (A⁺2). In addition, old sentences may acquire new support only as a result of discovered derivations from ϕ ((A⁺3)). (A⁺4) makes the simplifying assumption that adding ϕ does not disturb the preference relations already established; ϕ simply gets added in *some* appropriate position in the \preceq -induced chain of equivalence classes (and, hence, the non-uniqueness of focused expansion). We have implemented focused expansion in SNePS as assertion with forward inference, provided that only the sentences in \mathcal{F} are considered for matching.

Definition 4. A *focused revision* with *focus set* \mathcal{F} of \mathbb{S} with ϕ is a belief state $\mathbb{S} \dot{+}^{\mathcal{F}} \phi = \langle \mathcal{K}_{\dot{+}\mathcal{F}\phi}, \mathcal{B}_{\dot{+}\mathcal{F}\phi}, \sigma_{\dot{+}\mathcal{F}\phi}, \preceq_{\dot{+}\mathcal{F}\phi} \rangle$, satisfying the following properties.

- (A⁺1) *Base inclusion:* $\mathcal{B}_{\dot{+}\mathcal{F}\phi} \subseteq \mathcal{B}_{+\mathcal{F}\phi}$.
- (A⁺2) *Inclusion:* $\mathcal{K}_{\dot{+}\mathcal{F}\phi} \subseteq \mathcal{K}_{+\mathcal{F}\phi}$.
- (A⁺3) *Lumping:* $\psi \in \mathcal{K}_{+\mathcal{F}\phi} \setminus \mathcal{K}_{\dot{+}\mathcal{F}\phi}$ if and only if, for every $s \in \sigma_{+\mathcal{F}\phi}(\psi)$, $s \not\subseteq \mathcal{B}_{\dot{+}\mathcal{F}\phi}$.
- (A⁺4) *Preferential core-retainment:* $\psi \in \mathcal{B}_{+\mathcal{F}\phi} \setminus \mathcal{B}_{\dot{+}\mathcal{F}\phi}$ if and only if there is $\chi \in \mathcal{L}$ such that $(\chi \wedge \neg\chi) \in \text{Cn}_R(\mathcal{F}, \phi)$ and there is $s \in \sigma_{+\mathcal{F}\phi}(\chi \wedge \neg\chi)$ such that ψ is a minimal element of s with respect to $\preceq_{+\mathcal{F}}$.
- (A⁺5) *Support update:* For every $\psi \in \mathcal{L}$, $\sigma_{\dot{+}\mathcal{F}\phi}(\psi) = \sigma_{+\mathcal{F}\phi}(\psi) \cap 2^{\mathcal{B}_{\dot{+}\mathcal{F}\phi}}$
- (A⁺6) *Order preservation:* $\preceq_{\dot{+}\mathcal{F}\phi}$ is the restriction of $\preceq_{+\mathcal{F}\phi}$ to $\mathcal{B}_{\dot{+}\mathcal{F}\phi}$.

Thus, focused revision is focused assertion with forward inference followed by some kind of focused consolidation. Consolidation may be implemented by removing least-preferred beliefs (as per \preceq) from each support set of a contradiction ($\chi \wedge \neg\chi$) in the inconsistent belief state resulting from expansion by ϕ . As a result of consolidation, some base beliefs might be retracted in case focused expansion with ϕ results in a contradiction. (A⁺1) captures this intuition. (A⁺3) makes sure that only sentences that are still supported are believable. (A⁺4) guarantees that base beliefs that are evicted to retain (explicit) ϕ -relevant consistency indeed must be evicted. In addition, if a choice is possible, base beliefs that are least preferred are chosen for eviction. Note that, according to the above definition, this selection strategy is *skeptical*; that is, if multiple least preferred beliefs

exist, all are evicted. This strategy, however, is only adopted here to simplify the exposition, and nothing relevant depends on it.

The strongest result we can state about post-revision consistency is the following. The proof is eliminated for space limitations.

Theorem 1. *Let $\phi \in \mathcal{L}$ and $\mathbb{S} = \langle \mathcal{K}, \mathcal{B}, \sigma, \preceq \rangle$ be a belief state with $\mathcal{F} \subseteq \mathcal{K}$. For every, $\chi \in \mathcal{L}$, if $(\chi \wedge \neg\chi) \in \text{Cn}_R(\mathcal{F}, \phi)$, then $(\chi \wedge \neg\chi) \notin \mathcal{K}_{\dagger\mathcal{F}\phi}$.*

As a simple corollary, it follows that if the belief set \mathcal{K} is not known to be inconsistent, then, following focused revision, it is not known to be inconsistent.

Corollary 1. *Let $\mathbb{S} = \langle \mathcal{K}, \mathcal{B}, \sigma, \preceq \rangle$ be a belief state with $\mathcal{F} \subseteq \mathcal{K}$. For every $\phi, \chi \in \mathcal{L}$, if $(\chi \wedge \neg\chi) \notin \mathcal{K}$, then $(\chi \wedge \neg\chi) \notin \mathcal{K}_{\dagger\mathcal{F}\phi}$.*

This is as far as focused revision can claim about any notion of consistency restoration. In particular, if $\chi \wedge \neg\chi \in \mathcal{F}$ (not to mention \mathcal{K}), it is not necessary that $\chi \wedge \neg\chi \notin \mathcal{K}_{\dagger\mathcal{F}\phi}$. This is because focused revision is only guaranteed to treat an inconsistency to which the revising belief ϕ is relevant. Nevertheless, an inconsistency not induced by ϕ might be accidentally removed in case revising by ϕ side-effects the eviction of base beliefs that support the inconsistency.

4 Origins of Focus

To model high-level perception by focused belief revision, we need to consider interpretations of the focus set \mathcal{F} that are suitable for perception. Ultimately, we believe the choice of \mathcal{F} to be an empirical matter that can only be settled in the psychology or in the AI lab. Nevertheless, in this section, we present some general guidelines for the selection of \mathcal{F} in the context of perception.

We believe that the selection of a suitable focus set should be based on (at least) three factors: (i) what is vital for the agent, (ii) what is relevant for the agent, and (iii) how much resources are available for perception-induced reasoning.

For every agent, there are certain things that it cannot afford to not notice. For example, an agent might believe that, whenever there is fire, it should leave the building. (Remember agent *Y*?) A focus set of such an agent must include beliefs that allow it to conclude the imminence of fire from the perception of signs of fire. Thus, for every agent, there will be a certain set of vital beliefs, defined by fiat, that should be included in any focus set.

At a given time, to decide whether some belief is relevant *for* the agent we need to be explicit about two aspects of being relevant: (i) relevant *to* what, and (ii) relevant in what sense. We take a belief to be relevant for the agent if and only if it is relevant *to what the agent believes it is doing*. We believe that, broadly conceived, the set of acts the agent believes it is performing determines the focus of its attention, and, hence, provide a starting point for deciding on the contents of the focus set. Two acts in this set are either (i) independent concurrent acts, (ii) one of them is performed *by* performing the other, or (iii) one of them is performed as a step in (or towards) performing the other.

Now, what does it mean for a belief to be relevant to what the agent is doing? Intuitively, such a belief would be *about* the act the agent is performing, or about objects, individuals, and states involved in the act. This notion of *aboutness*, however, is notoriously elusive. One possibility is to explicitly associate beliefs to topics in the spirit of [15]. Another is to automatically partition the belief set \mathcal{K} using, for example, the greedy algorithm of Amir and McIlraith [16]. It should be noted, however, that Amir and McIlraith's concerns are very different from ours, and that, hence, their approach might not be suited for the case of perception. In this paper, we adopt a simple syntactic indicator of relevance that we call *nth-degree term sharing*. This is a generalization of the *variable sharing principle* adopted by relevance logicians for propositional languages [4]. The degree of term-sharing between two sentences reflects (in reverse proportion) the degree of relevance between the corresponding beliefs. As it turns out, the notion of degrees of relevance provides a way to tune the construction of the focus set to the amount of resources the agent can spend in the process: the more the resources, the more the beliefs with lower degrees of relevance the agent can consider.

For any $\mathbb{S} = \langle \mathcal{K}, \mathcal{B}, \sigma, \preceq \rangle$, $\alpha(\mathbb{S})$ (or α when \mathbb{S} is obvious) is the set of all sentences in \mathcal{K} of the form $Holds(Prog(a), * NOW)$. A function $\gamma : 2^{\mathcal{L}} \times \mathcal{L} \rightarrow 2^{\mathcal{L}}$ is a **relevance filtering function** if $\gamma(A, \phi) \subseteq A$. If $\phi \in \mathcal{L}$, $\tau(\phi)$ is the set of all closed terms occurring in ϕ and $TS(\phi) = \{\psi \mid \psi \in \mathcal{K} \cup \{\phi\} \text{ and } \tau(\phi) \cap \tau(\psi) \neq \emptyset\}$.

Definition 5. *Let $n \in \mathbb{N}$ and let γ be a relevance filtering function. An n^{th} -degree term sharing function with filter γ is a function $\mathfrak{t}_\gamma^n : \mathcal{L} \rightarrow 2^{\mathcal{L}}$ defined as follows:*

$$\mathfrak{t}_\gamma^n(\phi) = \begin{cases} \{\phi\} & \text{if } n = 0 \\ \gamma(TS(\phi), \phi) & \text{if } n = 1 \\ \{\psi \mid \text{for some } \xi \in \mathfrak{t}_\gamma^{n-1}(\phi), \psi \in \mathfrak{t}_\gamma^1(\xi)\} & \text{otherwise} \end{cases}$$

$\mathfrak{t}_\gamma^1(\phi)$ is the result of filtering the set of sentences that share at least one term with ϕ . The filtering function is used to account for the fact that term sharing is not sufficient for relevance. By experimenting with some filtering functions, we have found the following heuristics useful:

1. *If ϕ is of the form $Holds(s, t)$, filter out sentences of the form $Holds(s', t)$, for $\tau(s') \cap \tau(s) = \emptyset$. The intuition is that merely having two states hold simultaneously does not mean that propositions asserting these facts are relevant to one another.*
2. *In a language where properties are denoted by terms (which is the case for SNePS-based systems), if ϕ attributes a property to an individual, filter out sentences that attribute the same property to an unrelated individual. For example, if the agent sees a dog, what is relevant is not which other individuals are dogs, but general properties of dogs that may be used in reasoning about the perceived instance.*

Nevertheless, the exact definition of the filtering function largely depends on pragmatic factors that vary from one agent to another. In what follows, $\mathbb{S} = \langle \mathcal{K}, \mathcal{B}, \sigma, \preceq \rangle$ and $p = P_m(s, * NOW) \in \mathcal{L}$.

Definition 6. A *focus structure* $\mathbb{F}_{\mathbb{S},p}$ is a quadruple $\langle \mathcal{V}, \Gamma, \Delta, \rho \rangle$, where

- $\mathcal{V} \subseteq \mathcal{K}$ is a set of *vital beliefs*,
- $\Gamma : \alpha \cup \{p\} \longrightarrow [2^{\mathcal{L}} \times \mathcal{L} \longrightarrow 2^{\mathcal{L}}]$ is a mapping that assigns to every member of $\alpha \cup \{p\}$ a *relevance filtering function*,
- $\Delta : \alpha \cup \{p\} \longrightarrow \mathbb{N}$ is a mapping that assigns to every member of $\alpha \cup \{p\}$ a *degree of relevance*, and
- $\rho : 2^{\mathcal{L}} \times 2^{2^{\mathcal{L}}} \longrightarrow 2^{\mathcal{L}}$ is a *relevance choice function*, where

$$\rho(\mathfrak{t}_{\Gamma(p)}^{\Delta(p)}(p), \{\mathfrak{t}_{\Gamma(a)}^{\Delta(a)}(a)\}_{a \in \alpha}) \subseteq \bigcup_{\phi \in \alpha \cup \{p\}} \mathfrak{t}_{\Gamma(\phi)}^{\Delta(\phi)}(\phi)$$

The above notion of focus structures is an attempt to pinpoint the factors contributing to the construction of focus sets. Nonetheless, the definition is flexible enough to accommodate agent-specific considerations regarding vital beliefs, filtering functions, degrees of relevance to be considered based on available resources, and the construction of the set of relevant beliefs based on a relevance choice function.

Definition 7. Let $\mathbb{F}_{\mathbb{S},p} = \langle \mathcal{V}, \Gamma, \Delta, \rho \rangle$ be a focus structure. The *high-level perception* of s in \mathbb{S} with focus structure $\mathbb{F}_{\mathbb{S},p}$ is the focused belief revision, $\mathbb{S} \dot{+}^{\mathcal{F}} p$, of \mathbb{S} with p where

$$\mathcal{F} = \mathcal{V} \cup \Pi \cup \rho(\mathfrak{t}_{\Gamma(p)}^{\Delta(p)}(p), \{\mathfrak{t}_{\Gamma(a)}^{\Delta(a)}(a)\}_{a \in \alpha})$$

and p is a maximal element of $\mathcal{B}_{+\mathcal{F}p}$ with respect to $\preceq_{+\mathcal{F}p}$.

The set Π appearing in the definition of \mathcal{F} above is the perception theory referred to in Section 2. The requirement that p be a maximally preferred belief reflects the idea (often discussed in the philosophical literature) that having a present perceptual experience of some state s is indefeasible. Given the perception-belief axiom, an agent might not accept the proposition that s currently holds, if it has reason to believe that its sensory apparatus is faulty or that it is having an illusion. Nevertheless, the agent would still believe that it did have a perceptual experience of s . It is reasonable to assume that beliefs in Π about perception being veridical are not maximal elements of $\preceq_{+\mathcal{F}p}$. If this is indeed the case, high-level perception will satisfy the AGM postulate of success.

5 Conclusion

We have presented a framework for high-level perception as focused belief revision. This simultaneously addresses two issues. On one hand, the defeasibility of perception-induced beliefs is accounted for through the underlying reason maintenance system. On the other hand, bounded reflection on the contents of perception is implemented by two aspects of our system. First, the use of relevance logic guarantees that all perception-induced beliefs follow from the perception belief itself. Second, the definition of the focus set limits reasoning only to what

is relevant and vital for the agent, while taking issues of resource boundedness into account in a fairly general way.

We have identified four factors that, we believe, determine the construction of the focus set: the set of vital beliefs; the relevance filtering functions; the degrees of relevance; and the relevance choice function. All these factors depend on the belief state, the percept, and whatever the agent is doing. In the future, more experiments with different focus structures and different agents are needed to further evaluate the proposed framework.

References

1. Hexmoor, H., Lammens, J., Shapiro, S.C.: Embodiment in GLAIR: a grounded layered architecture with integrated reasoning for autonomous agents. In: *Proceedings of The Sixth Florida AI Research Symposium (FLAIRS 1993)*, pp. 325–329 (1993)
2. Shapiro, S.C., Ismail, H.O.: Anchoring in a grounded layered architecture with integrated reasoning. *Robotics and Autonomous Systems* 43(2-3), 97–108 (2003)
3. Alchourron, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic* 50(2), 510–530 (1985)
4. Anderson, A., Belnap, N.: *Entailment*, vol. I. Princeton University Press, Princeton (1975)
5. Shoham, Y., del Val, A.: A logic for perception and belief. Technical Report STAN-CS-91-1391, Department of Computer Science, Stanford University (1991)
6. Musto, D., Konolige, K.: Reasoning about perception. In: *Papers from the 1993 AAAI Spring Symposium on Reasoning about Mental States*, pp. 90–95 (1993)
7. Bell, J., Huang, Z.: Seeing is believing. In: *Papers from the 4th Symposium on Logical Formalizations of Commonsense Reasoning, Commonsense 1998* (1998)
8. Wooldridge, M., Lomuscio, A.: A computationally grounded logic of visibility, perception, and knowledge. *Logic Journal of the IGPL* 9(2), 273–288 (2001)
9. Wassermann, R.: Resource bounded belief revision. *Erkenntnis* 50(2-3), 429–446 (1999)
10. Shapiro, S.C., Rapaport, W.J.: The SNePS family. *Computers and Mathematics with Applications* 23(2-5), 243–275 (1992)
11. Martins, J., Shapiro, S.C.: A model for belief revision. *Artificial Intelligence* 35(1), 25–79 (1988)
12. Bacchus, F., Halpern, J., Levesque, H.: Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence* 111(1-2), 171–208 (1999)
13. Ismail, H.O.: A reason maintenance perspective on relevant Ramsey conditionals. *Logic Journal of the IGPL* 18(4), 508–529 (2010)
14. Johnson, F.L., Shapiro, S.C.: Dependency-directed reconsideration: Belief base optimization for truth maintenance systems. In: *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI 2005)*, pp. 313–320 (2005)
15. Demolombe, R., Jones, A.J.I.: Reasoning about topics: Towards a formal theory. In: *Workshop on Formalizing contexts: Papers from the 1995, AAAI Fall Symposium, Technical Report FS-95-02, Boston*, 55–59 (1995)
16. Amir, E., McIlraith, S.: Partition-based logical reasoning for first-order and propositional theories. *Artificial Intelligence* 162(1-2), 49–88 (2005)