Sensor Fusion in the Epistemic Situation Calculus

Christoph Schwering and **Tim Niemueller** and **Gerhard Lakemeyer**[1]
Nichola Abdo and **Wolfram Burgard**[2]

Abstract. Robot sensors are usually subject to error. Since in many practical scenarios a probabilistic error model is not available, sensor readings are often dealt with in a hard-coded, heuristic fashion. In this paper, we propose a logic to address the problem from a KR perspective. In this logic the epistemic effect of sensing actions is deferred to so-called fusion actions, which may resolve discrepancies and inconsistencies of recent sensing results. Moreover, a local closed world assumption can be applied dynamically. When needed, this assumption can be revoked and fusions can be undone using a form of forgetting.

1 INTRODUCTION

Even for supposedly straightforward tasks a robot needs to perform complex perception to gather sufficient knowledge about the environment and objects. Imagine a robot the with the goal of retrieving a mug from a table. Now the mug is not necessarily the only object on the table. In this example, we assume there are also a coffee pot and a sugar pack (as shown in Figure 1). The sensor is an RGB-D camera which provides 3D position and color information for each pixel, typically represented as a point cloud. Inherently such sensors perceive only the parts of an object facing the camera. For example, the mug, which has a handle, might be mistaken for a cup (without a handle) from the perspective of the robot in Figure 1. But when facing the table from the short edge and combining the previous sensor readings with the new ones, the robot can be certain about the object type. In our scenario, the sugar pack is only visible from that location, because it is otherwise occluded by the coffee pot. This example shows that it is sometimes necessary for the robot to look from more than just one perspective to understand a scene.

The reasoning formalism for high-level control must capture these ambiguities to allow for a robot to reason about its current state of knowledge about the world. This leads to *active perception*, that is to automatically execute actions to acquire the information necessary for the actual goal, or to reach certainty that it is not available.

In this paper, we address the outlined problem in the situation calculus. We present the modal first-order logic $\mathcal{ESF}$, which is intended to deal with incorrect and incompatible sensing results. In $\mathcal{ESF}$, sensing actions have no immediate effect to avoid inconsistent knowledge. Instead, sensing results are memorized and then merged by dedicated sensor fusion actions. Furthermore, actions may enforce a local closed world assumption to solidify the agent's opinion on certain things. For example, after looking at the table from various perspectives, the robot could fuse these sensings. Then it believes[3] that some objects are on the table. After further closing the domain of objects on the table, it believes that nothing else is on the table. To undo the epistemic effects of actions, we incorporate a simple notion of forgetting into $\mathcal{ESF}$.

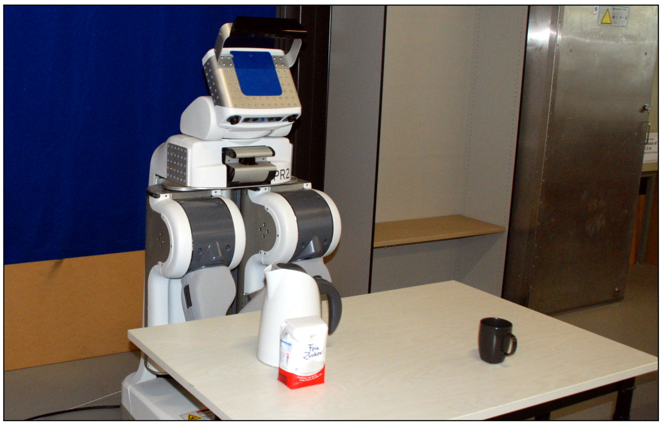

Figure 1: A PR2 robot looking at a table. From the current perspective, the sugar pack is occluded by the coffee pot and the handle of the mug is not visible, causing the robot to confuse it with a cup. From a different perspective, however, the robot would see the sugar pack and recognize the mug correctly. That is, sensings from different perspectives are inconsistent.

The paper is organized as follows. In the next section we present the logic $\mathcal{ESF}$ and show a few properties. In Section 3 we model two different scenarios with $\mathcal{ESF}$. While the first example is meant to familiarize with $\mathcal{ESF}$, the second one discusses our motivating tabletop scenario. After discussing related work in Section 4, we conclude.

2 THE LOGIC $\mathcal{ESF}$

$\mathcal{ESF}$ is a first-order modal logic for reasoning about actions and knowledge. It is a variant of the logic $\mathcal{ES}$ proposed by Lakemeyer and Levesque in [7]. While they recently proposed an extended version in [8], we refer to the original logic to simplify the presentation.

2.1 The Language

The language $\mathcal{ESF}$ consists of *fluent predicates* and *rigid terms*. The set of terms is the least set such that

- every first-order variable is a term;
- if f is a k-ary function symbol and $t_1, \ldots, t_k$ are terms, $f(t_1, \ldots, t_k)$ is a term.

[1] Knowledge-Based Systems Group, RWTH Aachen University, Germany, email: {schwering, niemueller, lakemeyer}@kbsg.rwth-aachen.de
[2] Autonomous Intelligent Systems, University of Freiburg, Germany, email:{abdon, burgard}@informatik.uni-freiburg.de

[3] In this paper we use the terms knowledge and belief interchangeably.

The set of well-formed formulas is the least set such that

- if P is a k-ary predicate symbol and $t_1, \ldots, t_k$ are terms, then $P(t_1, \ldots, t_k)$ is an (atomic) formula;
- if t_1, t_2 are terms, then $(t_1 = t_2)$ is a formula;
- if α and β are formulas and x is a variable, then $(\alpha \wedge \beta)$, $\neg\alpha$, $\forall x.\alpha$ are formulas;
- if α is a formula and t is a term, $[t]\alpha$ and $\Box\alpha$ are formulas;
- if α is a formula, n a natural number, and t a term, then $\mathbf{S}_t^n\alpha$, $\mathbf{K}\alpha$, $\mathbf{O}\alpha$ are formulas.

We read $[t]\alpha$ as "α holds after action t" and $\Box\alpha$ as "α holds after any sequence of actions." $\mathbf{S}_t^n\alpha$ is read as "α was sensed by the nth from last occurrence of action t," $\mathbf{K}\alpha$ as "α is known," and $\mathbf{O}\alpha$ as "α is all that is known."

We will use $\vee$, $\exists$, $\supset$, $\equiv$, *False*, *True* as the usual abbreviations. We let $\mathbf{K_{if}}\alpha$ stand for $\mathbf{K}\alpha \vee \mathbf{K}\neg\alpha$, which is read as "it is known whether or not α." We often omit universal quantifiers with maximum scope. When we omit brackets, connectives are ordered by increasing precedence: $\Box$, $\forall$, $\exists$, $\equiv$, $\supset$, $\vee$, $\wedge$, $\mathbf{K_{if}}$, $\mathbf{K}$, $\mathbf{O}$, $\mathbf{S}_t^n$, $[t]$, $\neg$.

Instead of having different sorts of objects and actions, we lump both sorts together and allow ourselves to use any term as an action or as an object. There are three distinguished predicates:

- *Poss*(a) expresses that action a is executable;
- *SR*(a, x) holds if x is a sensing result of a;
- *CW*(a, x) represents the closed world assumption made by a.

We call a formula without free variables a *sentence*. α_t^x denotes the result of substituting t for the free variable x in α. A formula with no $[t]$, $\Box$, $\mathbf{S}_t^n$, $\mathbf{K}$, or $\mathbf{O}$ is called a *fluent* formula. A formula with a single free variable a and with no $[t]$, $\Box$, $\mathbf{K}$, or $\mathbf{O}$ is called a *fusion* formula. We denote fusion formulas by the letter ϕ.

2.2 The Semantics

Truth of a sentence α after an action sequence z in $\mathcal{ESF}$ is defined wrt a (real) world w, a set of possible worlds e, a sensing history h, and a fusion formula ϕ.[4] We write $\phi, e, w, h, z \models \alpha$ for truth of α. The set of possible worlds e is called the epistemic state. In an objective context, w is indeed the real world, which is replaced with a $w' \in e$ in a subjective context. A world is a function from the ground atomic sentences and ground sequences of actions to $\{0, 1\}$. Let R denote the set of ground terms and R^* the set of sequences of ground terms, including the empty sequence $\langle\rangle$. Thus R can be considered the fixed domain of discourse of $\mathcal{ESF}$. A sensing history h maps each non-empty action sequence $z \cdot r$ to a set of worlds $h(z, r)$, which represents the worlds compatible with the sensing result of r after z. Finally ϕ is a fusion formula which asserts that the current world agrees on the fusion result of action a, the only free variable in ϕ. When writing $\phi, e, w, h, z \models \alpha$, we may leave out parts of the model irrelevant to α. For example, we often omit z and/or h if z is $\langle\rangle$, and if α is a fluent sentence, we omit ϕ, e, and h.

Each action r may in principle yield countably infinitely many sensing results, namely those terms s such that *SR*(r, s) holds. The epistemic state sensed by a ground action r after actions z is the set of worlds compatible with the sensing results purported by the real world w. We write $h_r^{w,z}$ for the sensing history h updated by the sensing results of r after z, which is defined by $h_r^{w,z}(z', r') = h(z', r')$ for all $z' \cdot r' \neq z \cdot r$ and by:

$$h_r^{w,z}(z, r) = \{w' \mid w'[SR(r, s), z] = w[SR(r, s), z] \text{ for all } s \in R\}$$

We use $h_{z'}^{w,z}$ as a shorthand for h updated with the sensing results throughout $z' \in R^*$. We define $h_0(z, r)$ to be the set of all worlds for all z and r.

Sensing results of r do not affect knowledge before they are fused by some action r'. The idea is that r' may fuse the (possibly contradicting) results of the n latest occurrences of r. We write $|z|$ for the length of a sequence z, and $|z|_r$ for the number of occurrences of r in z. We define $z|_r^n$ to be the longest prefix of z which does not contain the most recent n occurrences of action r. For example, $\langle r, s, r, s, r\rangle|_r^2 = \langle r, s\rangle$. (In case that $|z|_r < n$, the expression $|z|_r^n$ is undefined.)

Now we are ready to define the objective semantics:

1. $\phi, e, w, h, z \models P(r_1, \ldots, r_m)$ iff $w[P(r_1, \ldots, r_m), z] = 1$
2. $\phi, e, w, h, z \models (r = s)$ iff r and s are identical
3. $\phi, e, w, h, z \models (\alpha \wedge \beta)$ iff $\phi, e, w, h, z \models \alpha$ and $\phi, e, w, h, z \models \beta$
4. $\phi, e, w, h, z \models \neg\alpha$ iff $\phi, e, w, h, z \not\models \alpha$
5. $\phi, e, w, h, z \models \forall x.\alpha$ iff $\phi, e, w, h, z \models \alpha_r^x$ for all $r \in R$
6. $\phi, e, w, h, z \models [r]\alpha$ iff $\phi, e, w, h_r^{w,z}, z \cdot r \models \alpha$
7. $\phi, e, w, h, z \models \Box\alpha$ iff $\phi, e, w, h_{z'}^{w,z}, z \cdot z' \models \alpha$ for all $z' \in R^*$
8. $\phi, e, w, h, z \models \mathbf{S}_r^n\alpha$ iff $|z|_r \geq n$ and for all w', if $w' \in e \cap h(z|_r^n, r)$, then $\phi, e, w', h, z \models \alpha$

Notice that $[t]$ and $\Box$ update the sensing history h accordingly. The sensing history is used in rule 8, which defines $\mathbf{S}_r^n\alpha$ as truth of α in the epistemic state of the nth last occurrence of r. Observe that the sensed epistemic state is intersected with e and α is evaluated wrt z. This guarantees that sensing results are adequately projected into the current situation. Note that $\mathbf{S}_r^0\alpha$ refers to the epistemic state sensed by r in the current situation. If there is no nth last occurrence of r yet, $\mathbf{S}_r^n$*True* is false.

To characterize what is known when a sequence of actions z with sensing history h and fusion formula ϕ is performed in an epistemic state e, we define $e{\downarrow}^{\phi,h,z}$. Intuitively, $e{\downarrow}^{\phi,h,z}$ retains those worlds from e which agree with the sensing results from h after fusion according to ϕ and z and are compatible with the local CWA throughout z.[5] These conditions translate to the following definition:

- $w' \in e{\downarrow}^{\phi,h,\langle\rangle}$ iff $w' \in e$;
- $w' \in e{\downarrow}^{\phi,h,z\cdot r}$ iff
  (a) $w' \in e{\downarrow}^{\phi,h,z}$,
  (b) $e, w', h, z \models \phi_r^a$, and
  (c) for all $s \in R$, for all $w_1, w_2 \in e{\downarrow}^{\phi,h,z}$, if $w_1[CW(r, s), z] \neq w_2[CW(r, s), z]$, then $w'[CW(r, s), z] = 0$.

For an example of a fusion formula, suppose a fluent predicate $D(d)$ expresses that a robot's distance to a wall currently is d. The robot senses the distance through the action s and the action f shall fuse the latest two sensings by taking their average. The corresponding fusion formula is $a = \mathsf{f} \supset \exists d_1, d_2.\mathbf{S}_\mathsf{s}^1 D(d_1) \wedge \mathbf{S}_\mathsf{s}^2 D(d_2) \wedge D(\frac{d_1+d_2}{2})$. Then only those worlds are kept in $e{\downarrow}^{\phi,h,z\cdot\mathsf{f}}$ where the robot's distance is the average of the latest sensed distances. For an example of a CWA, suppose a robot looks at a table for objects and *On*(o) holds iff object o is on the table. After inspecting the table from every angle, the robot may want to assume it has seen every object on it. This is reflected by an action c which has no physical effect but makes a CWA on *On*(o) through *CW*$(\mathsf{c}, o) \equiv$ *On*(o). We give more examples in Section 3.

[4] The fusion formula ϕ is part of the model for technical reasons. Unlike $\mathcal{ES}$ [8] and the Scherl-Levesque [15] framework, we cannot use a single predicate and thus keep ϕ in the theory, because truth of ϕ usually does not depend on a single world but also on the sensing history, which is subject to change over the course of action and in introspective contexts.

[5] $\mathcal{ES}$ uses a relation $\simeq_z$ for a similar purpose.

The subjective semantics follows:

9. $\phi, e, w, h, z \models \mathbf{K}\alpha$ iff for all w',
   if $w' \in e{\downarrow}^{\phi,h,z}$, then $\phi, e, w', h, z \models \alpha$
10. $\phi, e, w, h, z \models \mathbf{O}\alpha$ iff for all w',
   $w' \in e{\downarrow}^{\phi,h,z}$ iff $\phi, e, w', h, z \models \alpha$

A set of sentences Σ entails α wrt a fusion formula ϕ (written $\Sigma \models_\phi \alpha$) iff for all e and w, if $\phi, e, w, h_0, \langle\rangle \models \sigma$ for all $\sigma \in \Sigma$, then $\phi, e, w, h_0, \langle\rangle \models \alpha$. A set of sentences Σ entails a sentence α (written $\Sigma \models \alpha$) iff for all fusion formulas ϕ, $\Sigma \models_\phi \alpha$. A sentence α is valid wrt ϕ (written $\models_\phi \alpha$) iff $\{\} \models_\phi \alpha$. A sentence is valid (written $\models \alpha$) iff $\{\} \models \alpha$.

For the rest of the subsection, we investigate introspection and the local CWA in $\mathcal{ESF}$. To begin with, $\mathcal{ESF}$ is fully introspective:

Theorem 1 $\models \Box\mathbf{K}\alpha \supset \mathbf{K}\mathbf{K}\alpha$ *and* $\models \Box\neg\mathbf{K}\alpha \supset \mathbf{K}\neg\mathbf{K}\alpha$.

Proof. Both properties hold because $e{\downarrow}^{\phi,h,z}$ for $(\neg)\mathbf{K}\alpha$ is the same as for $\mathbf{K}(\neg)\mathbf{K}\alpha$. □

The following theorem expresses that after executing an action r, for each s either $CW(r, s)$ is known or $\neg CW(r, s)$ is known. For the aforementioned example this means that after action c, we know for each object s whether or not $On(s)$ is true, that is, if s is on the table or not. Proviso for the theorem is that r has no physical effect on truth of $CW(r, s)$:

Theorem 2 $\models \Box\mathbf{K}([a]CW(a, x) \equiv CW(a, x)) \supset [a]\mathbf{K}_{\mathbf{if}}CW(a, x)$.

Proof. Let $\phi, e, w, h^w_z, z \models \mathbf{K}([r]CW(r, s) \equiv CW(r, s))$.

If $\phi, e, w, h^w_z, z \models \mathbf{K}CW(r, s)$, condition (c) in $e{\downarrow}^{\phi,h^w_{z\cdot r},z\cdot r}$ has no effect. Since $e{\downarrow}^{\phi,h^w_{z\cdot r},z\cdot r} \subseteq e{\downarrow}^{\phi,h^w_z,z}$ and r does not change $CW(r, s)$, it follows $\phi, e, w, h^w_{z\cdot r}, z \cdot r \models \mathbf{K}CW(r, s)$.

Now suppose $\phi, e, w, h^w_z, z \models \neg\mathbf{K}CW(r, s)$. Then condition (c) requires $w'[CW(r, s), z] = 0$ for all $w' \in e{\downarrow}^{\phi,h^w_{z\cdot r},z\cdot r}$. As r does not change $CW(r, s)$, it follows $\phi, e, w, h^w_{z\cdot r}, z \cdot r \models \mathbf{K}\neg CW(r, s)$. □

2.3 Basic Action Theories

We define the $\mathcal{ESF}$ variant of Reiter's basic action theories [14]. A basic action theory (BAT) over a finite set of fluents $\mathcal{F}$ contains sentences which describe the initial situation, action preconditions, and both the actions' physical and epistemic effects. In $\mathcal{ESF}$ we distinguish between the objective BAT Σ and the BAT Σ' subjectively known to the agent:

$$\Sigma = \Sigma_0 \cup \Sigma_{pre} \cup \Sigma_{post} \cup \Sigma_{sense} \text{ and}$$
$$\Sigma' = \Sigma'_0 \cup \Sigma_{pre} \cup \Sigma_{post} \cup \Sigma'_{sense} \cup \Sigma'_{close}$$

The components of Σ and Σ' are as follows:[6]

- Σ_0 is a set of fluent sentences which hold initially;
- Σ'_0 is a set of fluent sentences the agent believes to be true;
- Σ_{pre} is a singleton sentence of the form $\Box Poss(a) \equiv \alpha$;
- Σ_{post} contains for every $F \in \mathcal{F}$ a sentence $\Box[a]F(\vec{x}) \equiv \alpha$;
- Σ_{sense} is a singleton sentence of the form $\Box SR(a, s) \supset \alpha$;
- Σ'_{sense} is a singleton sentence of the form $\Box SR(a, s) \equiv \alpha$;
- Σ'_{close} is a singleton sentence of the form $\Box CW(a, x) \equiv \alpha$;

[6] We abuse notation and do not distinguish finite sets of sentences from conjunctions.

where all α's are fluent formulas. The sentences in Σ_{post} are called *successor state axioms* (SSAs). They follow the simple pattern $\Box[a]F(\vec{x}) \equiv \gamma^+_F(a, \vec{x}) \vee F(\vec{x}) \wedge \neg\gamma^-_F(a, \vec{x})$ where $\gamma^\pm_F(a, \vec{x})$ capture the positive and negative effects of a on $F(\vec{x})$. This pattern is key to Reiter's solution to the frame problem [14]. Observe that Σ and Σ' not only differ in Σ_0 and Σ'_0 (as is common in $\mathcal{ES}$), but also in Σ_{sense} and Σ'_{sense}. The idea is that Σ_{sense} merely constrains the possible sensing results to rule out implausible values.

Intuitively, the fusion formula ϕ should be part of the BATs, too, as it is intended to fuse sensing results. For technical reasons, ϕ must be a parameter of the semantics, though.

For example BATs we refer to Section 3.

2.4 Forgetting

In many settings it may be desirable to forget some information, for example, to revoke a sensor fusion or a local CWA. Rajaratnam et al. [12] proposed an extension of the Scherl-Levesque approach [15] to knowledge in the situation calculus. We adopt their ideas for our logic $\mathcal{ESF}$ in this subsection.

The action $\mathsf{forget}(r)$ undoes the epistemic effect of the last occurrence of r. We replace condition (a) of $e{\downarrow}^{\phi,h,z\cdot r}$ with

(a') $w' \in e{\downarrow}^{\phi,h,z'}$ where z' is z with the last occurrence of s removed if $r = \mathsf{forget}(s)$ for some s, otherwise z' is just z.

We show that if r has no physical effect on α but through its epistemic effect (may it be sensing or a CWA) tells the agent that α holds, then this awareness of α can be revoked with the action $\mathsf{forget}(r)$. Of course, this only holds if $\mathsf{forget}(r)$ itself has no physical effect on α and does not itself remove any possible worlds through sensing or a CWA. This proviso is expressed by β in the following theorem:

Theorem 3 *Let ϕ be a fusion formula and β stand for* $([a]\alpha \equiv \alpha) \wedge [a]([\mathsf{forget}(a)]\alpha \equiv \alpha) \wedge [a]\phi^a_{\mathsf{forget}(a)} \wedge [a]\forall x. CW(\mathsf{forget}(a), x)$. *Then* $\models_\phi \Box\mathbf{K}\beta \wedge \neg\mathbf{K}\alpha \wedge [a]\mathbf{K}\alpha \supset [a][\mathsf{forget}(a)]\neg\mathbf{K}\alpha$.

Proof. Suppose $\phi, e, w, h^w_z, z \models \mathbf{K}\beta^a_r$. Then $\phi, e, w', h^w_z, z \models [r][\mathsf{forget}(r)]\alpha \equiv \alpha$ for all $w' \in e{\downarrow}^{\phi,h^w_z,z}$. Thus we only need to show that $e{\downarrow}^{\phi,h^w_{z\cdot r\cdot \mathsf{forget}(r)},z\cdot r\cdot \mathsf{forget}(r)} = e{\downarrow}^{\phi,h^w_z,z}$. This holds because condition (a') skips action conditions (b) and (c) for r, and conditions (b) and (c) are satisfied for $\mathsf{forget}(r)$ trivially for all $w' \in e{\downarrow}^{\phi,h^w_z,z}$. □

This is a simplified version of a theorem from [12]. It is straightforward to generalize to more actions between r and $\mathsf{forget}(r)$.

We remark that our definition is simpler than the one in [12]. The reason is that Rajaratnam et al. cannot drop actions from situations terms without confusing their accessibility relations.

3 EXAMPLES

In this section we model two scenarios as BATs and show a few properties. To begin with, we apply $\mathcal{ESF}$ to a variant of the running example of [8]. This example shows how the fused sensing result can differ from the sensings to be fused. More precisely, we will fuse two sensings by their disjunction, which would not be possible in the Scherl-Levesque framework and $\mathcal{ES}$ due to their monotonic nature [15, 8]. Afterwards, we model a tabletop perception scene. There we show how to deal with an unknown number of objects and exemplify use of the CWA.

3.1 Distance to the Wall

Imagine a robot moving towards a wall. The robot's initial distance to the wall is 5 units (written as $D(5)$), but it does not know this fact. By Γ we denote an axiomatization of the rational numbers, which we need to work with distances. Thus we have for the initial situation:

$$\Sigma_0 \doteq \{D(d) \equiv d = 5\} \cup \Gamma$$
$$\Sigma'_0 \doteq \{\exists d.D(d),\ \forall d, d'.D(d) \wedge D(d') \supset d = d'\} \cup \Gamma$$

The robot may move one unit towards the wall (through action m).[7] The appropriate precondition axiom and SSA are:

$$\Sigma_{pre} \doteq \{\Box Poss(a) \equiv (a = \mathsf{m} \supset \neg D(0))\}$$
$$\Sigma_{post} \doteq \{\Box D(d) \equiv \exists d'.a = \mathsf{m} \wedge D(d') \wedge d = d' - 1 \vee D(d) \wedge a \neq \mathsf{m}\}$$

The robot is equipped with a sonar sensor (action s), which yields intervals of possible distances. This is captured by Σ_{sense}, while Σ'_{sense} expresses which possible worlds agree on a specific distance:

$$\Sigma_{sense} \doteq \{\Box SR(a, x) \supset (a = \mathsf{s} \supset \exists d_1, d_2.x = [d_1, d_2])\}$$
$$\Sigma'_{sense} \doteq \{\Box SR(a, x) \equiv (a = \mathsf{s} \supset \exists d_1, d_2.x = [d_1, d_2] \wedge \exists d.D(d) \wedge d_1 \leq d \leq d_2)\}$$

Notice that a sensing result $[d_1, d_2]$ represents disjunctive information. Disjunctive information must be encoded within such a single sensing result, as opposed to a set of sensing results like $SR(\mathsf{s}, d)$ for all $d_1 \leq d \leq d_2$. This is because the set of sensing results is interpreted conjunctively: a world is compatible with a sensing only if it agrees with all sensing results.

Since the robot mistrusts its own sensor, it takes the fusion (action f) of two sensings to be the union of the reported intervals. That is, a possible world's distance must be considered possible in one of the last sensings. Thus we use as fusion formula

$$\phi \doteq a = \mathsf{f} \supset \exists d.D(d) \wedge (\neg\mathbf{S}^1_{\mathsf{s}}\neg D(d) \vee \neg\mathbf{S}^2_{\mathsf{s}}\neg D(d))$$

Notice that $\neg\mathbf{S}^1_{\mathsf{s}}\neg D(d)$ expresses that d was not ruled out by the last sensing of s.

We do not use the local CWA here, so we define:

$$\Sigma'_{close} \doteq \{\Box CW(a, x) \equiv True\}$$

Now we can reason about what is entailed by this BAT. Let e, w be such that $\phi, e, w \models \Sigma \wedge \mathbf{O}\Sigma'$ and $w \models SR(\mathsf{s}, [4, 7]) \wedge [\mathsf{s}]SR(\mathsf{s}, [3, 6])$, that is, the first s reports the interval $[4, 7]$ and a subsequent s reports $[3, 6]$. We show the following properties:

1. $\phi, e, w, h \models [\mathsf{s}][\mathsf{s}][\mathsf{f}]\mathbf{K}(D(d) \equiv 3 \leq d \leq 7)$
   The robot believes the distance is in $[3, 7]$:
   Let g stand for $h^w_{\langle \mathsf{s},\mathsf{s},\mathsf{f}\rangle}$ and let $w' \in e\downarrow^{\phi,g,\langle \mathsf{s},\mathsf{s},\mathsf{f}\rangle}$. Then $e, w', g, \langle \mathsf{s}, \mathsf{s}\rangle \models \phi^a_{\mathsf{f}}$ due to condition (b) for $e\downarrow^{\phi,g,\langle \mathsf{s},\mathsf{s},\mathsf{f}\rangle}$, that is, $e, w', g, \langle \mathsf{s}, \mathsf{s}\rangle \models \exists d.D(d) \wedge (\neg\mathbf{S}^1_{\mathsf{s}}\neg D(d) \vee \neg\mathbf{S}^2_{\mathsf{s}}\neg D(d))$. We have $e, g, \langle \mathsf{s}, \mathsf{s}\rangle \models \neg\mathbf{S}^2_{\mathsf{s}}\neg D(d)$ iff $w'[D(d), \langle \mathsf{s}, \mathsf{s}\rangle] = 1$ for some $w' \in e \cap \{w' \mid w'[SR(\mathsf{s}, r), \langle\rangle] = w[SR(\mathsf{s}, r), \langle\rangle]$ for all $r \in R\}$ iff $4 \leq d \leq 7$. Analogously $e, g, \langle \mathsf{s}, \mathsf{s}\rangle \models \neg\mathbf{S}^1_{\mathsf{s}}\neg D(d)$ iff $3 \leq d \leq 6$. Thus the property follows.
2. $\phi, e, w, h \models [\mathsf{s}][\mathsf{s}][\mathsf{m}][\mathsf{f}]\mathbf{K}(D(d) \equiv 2 \leq d \leq 6)$
   Fusion projects the sensed distances by one unit towards the wall due to the m action:
   The proof is analogous to the previous one except that $\langle \mathsf{s}, \mathsf{s}\rangle$ is replaced with $\langle \mathsf{s}, \mathsf{s}, \mathsf{m}\rangle$, which leads to intervals $[2, 5]$ and $[3, 6]$ instead of $[3, 6]$ and $[4, 7]$.

[7] We use different typefaces to distinguish action function symbols from other terms: action vs *term*.

3.2 Tabletop Object Search

In our previous work [11], we presented a system for active perception where a robot navigates around a table in order to detect specific objects on it. Our approach highlighted the importance of merging sensor data from multiple perspectives to overcome problems like occlusions. In the current system, the robot perceives point clouds of the table scene with its Kinect camera. From this it extracts object clusters which are each assigned a unique object ID. These IDs remain stable among different perspectives. Additionally, it matches features extracted from available 3D object models to those computed from the depth image of the scene. Both kinds of observations are then combined, yielding a – possibly empty – type distribution for each object. The type detection, however, highly depends on the perspective of the camera. For example, the robot cannot necessarily distinguish a mug (with a handle) from a cup (without a handle) if the handle is not visible from the current perspective. This ambiguity must be resolved by observing the scene from multiple perspectives.

We now model this scenario in $\mathcal{ESF}$. The aforementioned stable object IDs allow us to use rigid terms to refer to the same object in different situations. We use the predicate $On(o)$ to express that object o is on the table and $Is(o, t)$ to say that o is of type t. For example, an object might have the type *Mug* (with handle) or *Cup* (without handle). To simplify matters, we only consider two perspectives: The robot either stands at the long (L) or the short side of the table ($\neg L$) and it can move from either position to the other (through action m). The robot may look on the table (action s) to see some objects and possibly recognize their type. Note that we do not deal with confidence values for type hypotheses in this work. Lastly there is an action to solidify the robot's view on what is on the table by enforcing a local CWA (c).

We proceed to define the objective and subjective BATs Σ and Σ'. Initially the robot is located at the long side and is aware of this fact. For the sake of simplicity in this example we further axiomatize that any object has exactly one type:

$$\Sigma_0 \doteq \Sigma'_0 \doteq \{L,\ \exists t.Is(o, t),\ Is(o, t) \wedge Is(o, t') \supset t = t'\}$$

There are no specific preconditions in this scenario:

$$\Sigma_{pre} \doteq \{\Box Poss(a) \equiv True\}$$

Only the physical value of L may change due to actions:

$$\Sigma_{post} \doteq \{\Box[a]L \equiv (a = \mathsf{m} \wedge \neg L) \vee (L \wedge a \neq \mathsf{m}),\ \Box[a]Is(o, t) \equiv Is(o, t),\ \Box[a]On(o) \equiv On(o)\}$$

Now we turn to Σ_{sense} and Σ'_{sense}. We know that the sensor only reports types, so we constrain the reported sensor values accordingly:

$$\Sigma_{sense} \doteq \{\Box SR(a, x) \supset (a = \mathsf{s} \supset \exists o.x = obj(o) \vee \exists o, t.x = type(o, t))\}$$

The subjective *SR* axioms shall express when a possible world agrees with sensing results:

$$\Sigma'_{sense} \doteq \{\Box SR(a, x) \equiv (a = \mathsf{s} \supset \exists o.x = obj(o) \wedge On(o) \vee \exists o, t.x = type(o, t) \wedge Is(o, t))\}$$

We use the following fusion scheme:

- If an object was seen on the table in either of the last two sensings, then the robot believes that it is on the table. That is, all worlds where that object is not on the table are considered impossible.

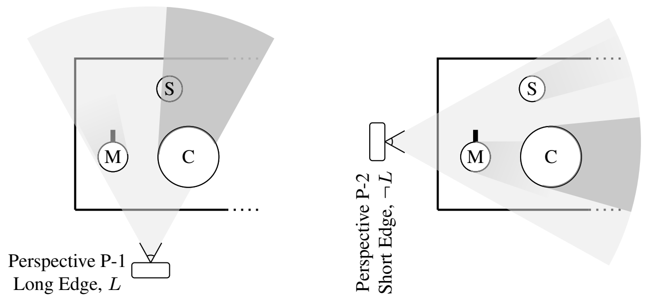

Figure 2: Tabletop with a mug M, a sugar box S, and a coffee pot C from two different perspectives. Light gray cones denote horizontal viewing angle, dark gray regions represent sensor shadows. From the first perspective, S is partly occluded by C. M's handle is only visible from the second perspective.

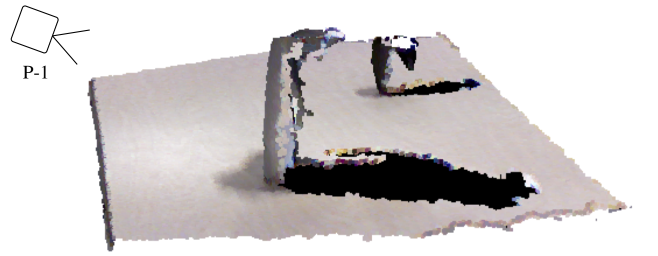

Figure 3: The point cloud as seen by the robot in Figure 1 and in the first perspective in Figure 2. The dark edge is the long edge of the table faced by the robot. Black color indicates shadow areas. Note that the sugar pack and mug handle are not visible.

- If an object was recognized as a t in either sensing, we believe it is a t, modulo one constraint: if t is *Cup*, then it must not be recognized as *Mug* in the other sensing. The idea behind this constraint is that often a *Mug* is mistaken for a *Cup* because its handle is not visible. In other words, sensing results *Mug* override sensing results *Cup*.

We translate this scheme to $\mathcal{ESF}$ formulas:

$$\alpha \doteq \forall o.\mathbf{S}^1_{\mathsf{s}} On(o) \vee \mathbf{S}^2_{\mathsf{s}} On(o) \supset On(o)$$

$$\beta \doteq \forall o,t.(\mathbf{S}^1_{\mathsf{s}} Is(o,t) \wedge (t = Cup \supset \neg \mathbf{S}^2_{\mathsf{s}} Is(o, Mug)) \vee \mathbf{S}^2_{\mathsf{s}} Is(o,t) \wedge (t = Cup \supset \neg \mathbf{S}^1_{\mathsf{s}} Is(o, Mug))) \supset Is(o,t)$$

Then we define our fusion formula as

$$\phi \doteq a = \mathsf{f} \supset \alpha \wedge \beta$$

The action c shall have the effect that the robot believes it has seen everything on the table, therefore we define:

$$\Sigma'_{close} \doteq \{\Box CW(a,x) \equiv (a = \mathsf{c} \supset On(x))\}$$

Suppose $\phi, e, w, h \models \Sigma \wedge \mathbf{O}\Sigma'$ for the remainder of this subsection. Imagine that the real world w is as depicted in Figure 2 with three objects M, C, S on the table, where M is a mug (with a handle), C is a large coffee pot, and S is a sugar box:

$$w \models On(M) \wedge On(C) \wedge On(S) \wedge Is(M, Mug) \wedge Is(C, Coffee) \wedge Is(S, Sugar)$$

However, the robot correctly identifies the mug only when positioned at the table's short side, otherwise it does not see the mug's handle and thus mistakes it for a cup. Furthermore the sugar box is hidden by the coffee pot when the robot is standing on the long edge. Except for the case when the robot does not see the mug's handle, the sensed position is correct. In logic:

$$\begin{aligned} w \models \Box SR(\mathsf{s}, x) \equiv x = obj(M) \vee \\ x = obj(C) \vee \\ x = type(C, Coffee) \vee \\ (L \wedge x = type(M, Cup)) \vee \\ (\neg L \wedge (x = obj(S) \vee \\ x = type(M, Mug) \vee \\ x = type(S, Sugar))) \end{aligned}$$

In this scenario, the following properties hold:

1. $\phi, e, w, h \models [\mathsf{s}][\mathsf{f}]\mathbf{K} Is(M, Cup)$
After sensing only from the long side, the robot erroneously thinks M is a *Cup*:
Let g stand for $h^w_{\langle \mathsf{s},\mathsf{f}\rangle}$. We have $e, g, \langle \mathsf{s}\rangle \models \mathbf{S}^1_{\mathsf{s}} Is(M, Cup)$ and $e, g, \langle \mathsf{s}\rangle \not\models \mathbf{S}^2_{\mathsf{s}} Is(M, Mug)$ since there is just one sensing. This and the definition of ϕ^a_{f} give that we have $w'[Is(M, Cup), \langle \mathsf{s}\rangle] = 1$ for all $w' \in e{\downarrow}^{\phi,g,\langle \mathsf{s},\mathsf{f}\rangle}$.
2. $\phi, e, w, h \models [\mathsf{s}][\mathsf{f}][\mathsf{m}][\mathsf{s}][\mathsf{forget}(\mathsf{f})][\mathsf{f}]\mathbf{K} Is(M, Mug)$
After sensing from both sides and forgetting the first fusion, the robot correctly believes M is a *Mug*:
Observe that $e{\downarrow}^{\phi,g,\langle \mathsf{s},\mathsf{f},\mathsf{m},\mathsf{s},\mathsf{forget}(\mathsf{f}),\mathsf{f}\rangle} = e{\downarrow}^{\phi,h^w_{\langle \mathsf{s},\mathsf{m},\mathsf{s},\mathsf{f}\rangle},\langle \mathsf{s},\mathsf{m},\mathsf{s},\mathsf{f}\rangle}$ where g stands for $h^w_{\langle \mathsf{s},\mathsf{f},\mathsf{m},\mathsf{s},\mathsf{forget}(\mathsf{f}),\mathsf{f}\rangle}$, because condition (a') means that $\mathsf{forget}(\mathsf{f})$ undoes the epistemic effects of the first f. Furthermore we have $e, g, \langle \mathsf{s},\mathsf{m},\mathsf{s}\rangle \models \mathbf{S}^1_{\mathsf{s}} Is(M, Mug)$. Thus $w'[Is(M, Mug), \langle \mathsf{s}\rangle] = 1$ for all $w' \in e{\downarrow}^{\phi,g,\langle \mathsf{s},\mathsf{m},\mathsf{s},\mathsf{f}\rangle}$.
3. $\phi, e, w, h \models [\mathsf{s}][\mathsf{f}]\neg\mathbf{K}(x = M \vee x = C \equiv On(x))$
After sensing once and fusing, the robot does not believe that M and C are the only objects on the table:
There is some $w' \in e$ such that $w' \models (x = M \vee x = C \vee x = r \supset On(s)) \wedge Is(M, Cup) \wedge Is(C, Coffee)$ for an arbitrary $r \notin \{M, C\}$. Since w' agrees to the fused sensing results, $w' \in e{\downarrow}^{\phi,h^w_{\langle \mathsf{s},\mathsf{f}\rangle},\langle \mathsf{s},\mathsf{f}\rangle}$. Hence $(x = M \vee x = C \equiv On(x))$ is not known.
4. $\phi, e, w, h \models [\mathsf{s}][\mathsf{f}][\mathsf{c}]\mathbf{K}(x = M \vee x = C \equiv On(x))$
After additionally closing the domain, the robot thinks M and C are the only objects on the table:
Let g stand for $h^w_{\langle \mathsf{s},\mathsf{f},\mathsf{c}\rangle}$. Notice that for all $w' \in e{\downarrow}^{\phi,g,\langle \mathsf{s},\mathsf{f},\mathsf{c}\rangle}$, $w', \langle \mathsf{s},\mathsf{f},\mathsf{c}\rangle \models On(M) \wedge On(C)$ holds because otherwise condition (b) for $e{\downarrow}^{\phi,g,\langle \mathsf{s},\mathsf{f}\rangle}$ would be violated. To see that $w'[On(r), \langle \mathsf{s},\mathsf{f},\mathsf{c}\rangle] = 0$ for all $r \notin \{M, C\}$, suppose $w_1, w_2 \in e{\downarrow}^{\phi,g,\langle \mathsf{s},\mathsf{f}\rangle}$ and $w_1[On(r), \langle \mathsf{s},\mathsf{f}\rangle] \neq w_2[On(r), \langle \mathsf{s},\mathsf{f}\rangle]$. Such w_1 and w_2 exist as argued in the previous property. Then due to condition (c) of $e{\downarrow}^{\phi,g,\langle \mathsf{s},\mathsf{f},\mathsf{c}\rangle}$, $w'[CW(\mathsf{c}, r), \langle \mathsf{s},\mathsf{f}\rangle] = 0$, which implies that $w'[On(r), \langle \mathsf{s},\mathsf{f}\rangle] = 0$.
5. $\phi, e, w, h \models [\mathsf{s}][\mathsf{m}][\mathsf{s}][\mathsf{f}]\mathbf{K}(x = M \vee x = C \vee x = S \supset On(x))$
After inspecting the table from both sides and fusing these sensings, the robot believes that M, C, S are on the table:
Let g stand for $h^w_{\langle \mathsf{s},\mathsf{m},\mathsf{s},\mathsf{f}\rangle}$. For each $r \in \{M, C, S\}$ we have $e, g, \langle \mathsf{s},\mathsf{m},\mathsf{s}\rangle \models \mathbf{S}^1_{\mathsf{s}} On(r)$ and thus by condition (b), $w'[On(r), \langle \mathsf{s},\mathsf{m},\mathsf{s},\mathsf{f}\rangle] = 1$ for each $w' \in e{\downarrow}^{\phi,g,\langle \mathsf{s},\mathsf{m},\mathsf{s},\mathsf{f}\rangle}$.

4 RELATED WORK

Reiter's situation calculus in its original form [14] does not account for sensing actions and the agent's knowledge or belief. An epistemic extension [15] by Scherl and Levesque added a possible worlds semantics within classical first-order logic. Lakemeyer and Levesque [8] gave a semantic account of that in the modal first-order logic $\mathcal{ES}$, which itself is the basis of $\mathcal{ESF}$. In both, the original Scherl-Levesque framework and $\mathcal{ES}$, actions have binary sensing results and after such a sensing action, the agent knows the sensing result immediately.

Other action formalisms like SADL [6], the event calculus [5], and the fluent calculus [18] have or can be extended to have a notion of knowledge, too, but they do not address the problem of contradictory sensing results and their fusion.

An extension of the epistemic situation calculus by Bacchus et al. [1] incorporates Bayesian belief update. This requires an error model in the form of an action likelihood function that formalizes the gap between reality and the robot's mind. For example, it may express that the actual result of a sonar sensor is normally distributed around the real distance. Action likelihoods give rise to a probability distribution of the possible worlds. While this distribution is discrete in [1], Belle et al.'s variant [2] allows for continuous ones. In scenarios we have in mind for $\mathcal{ESF}$, however, such a precise error model is not known.

Our sensing histories are somewhat related to IndiGolog's [3] concept of histories. It differs, however, in that we interpret sensings as epistemic states and have a notion of sensor fusion, whereas IndiGolog assumes correct sensors and adds the their binary results to the theory during on-line execution.

Shapiro et al. [16] presented a theory for belief change in the situation calculus. They specify initial beliefs with a $\Rightarrow$ operator in the spirit of counterfactuals. Sensing results then trigger belief change. This concept could be applied to the cup vs mug problem: we believe an object is a cup until we see it has a handle, in which case we believe it is a mug. In [16] sensing results are assumed to be correct, though. We allow for incompatible sensing results by deferring their epistemic effect to a fusion action. A belief revision scheme to deal with contradictory fusion results remains future work.

The closed world assumption was introduced by Reiter [13]. Etzioni et al. [4] applied a local closed world assumption in a dynamic environment. They also account for loss of closed world information, which we in a way allow by forgetting.

The forgetting mechanism of $\mathcal{ESF}$ is essentially the same as the one proposed by Rajaratnam et al. [12]. This notion of forgetting is fundamentally different from Lin and Reiter's logical forgetting [10]. We refer to [12] for the details.

KNOWROB [17] is a recent example for a robotic knowledge processing system. It acts as a database providing virtual knowledge bases which can be queried from the task reasoner. It gathers information from various sources like ontological databases or sensors. Each sensor detection is stored as a new instance. Queries then aim at retrieving the latest information, rather than fusing information and dealing with inconsistencies explicitly. Active perception is not performed by the system itself, but relies on an executive to orchestrate the proper action sequence.

5 CONCLUSION AND FUTURE WORK

In this paper we presented a logic for reasoning about knowledge in the presence of actions which may yield incorrect and incompatible sensing results. This addresses a real need we encounter in our work with robotics. To deal with this problem, our logic differs in several ways from most previous approaches. Firstly, it allows to sense an unbounded number of objects and the purported sensing results to vary over the course of action, contradicting themselves and reality. We remark that this does not necessitate second-order logic. Reasoning can be done by inspecting specific real worlds which specify certain (incorrect) sensing results. Secondly, sensing actions do not affect knowledge immediately. Instead, this effect is deferred to a fusion action. Thirdly, actions can apply a local closed world assumption to solidify the agent's episteme. Finally, the forgetting mechanism allows to mitigate the epistemic effects of preceding actions.

The next step is to extend $\mathcal{ESF}$ to deal with uncertainties in sensing results. We also plan to deploy a decidable subset in the spirit of $\mathcal{ESL}$ [9] on our robots. Lastly, we aim to integrate a theory of belief so that incompatible sensor fusions do not make the agent's knowledge inconsistent but revise his beliefs appropriately. In the context of a larger project on hybrid reasoning[8] we plan to make use of the concepts of $\mathcal{ESF}$ for active perception.

ACKNOWLEDGEMENTS

This work was supported by the German National Science Foundation (DFG) research unit FOR 1513 on Hybrid Reasoning for Intelligent Systems.[8]

REFERENCES

[1] Fahiem Bacchus, Joseph Y. Halpern, and Hector J. Levesque, 'Reasoning about noisy sensors and effectors in the situation calculus', *Artificial Intelligence*, **111**(1–2), (1999).
[2] Vaishak Belle and Hector J. Levesque, 'Reasoning about continuous uncertainty in the situation calculus', in *Proc. IJCAI*, (2013).
[3] Giuseppe De Giacomo and Hector Levesque, 'An incremental interpreter for high-level programs with sensing', in *Logical foundations for cognitive agents: contributions in honor of Ray Reiter*, (1999).
[4] Oren Etzioni, Keith Golden, and Daniel Weld, 'Tractable closed world reasoning with updates.', in *Proc. KR*, (1994).
[5] Jeremy Forth and Murray Shanahan, 'Indirect and conditional sensing in the event calculus', in *Proc. ECAI*, (2004).
[6] Keith Golden and Daniel Weld, 'Representing sensing actions: The middle ground revisited', in *Proc. KR*, (1996).
[7] Gerhard Lakemeyer and Hector J. Levesque, 'Situations, si! situation terms, no!', in *Proc. KR*, (2004).
[8] Gerhard Lakemeyer and Hector J. Levesque, 'A semantic characterization of a useful fragment of the situation calculus with knowledge', *Artificial Intelligence*, **175**(1), (2011).
[9] Gerhard Lakemeyer and Hector J. Levesque, 'Decidable reasoning in a fragment of the epistemic situation calculus', in *Proc. KR*, (2014).
[10] Fangzhen Lin and Ray Reiter, 'Forget it!', in *Proc. AAAI Fall Symposium on Relevance*, (1994).
[11] Tim Niemueller, Nichola Abdo, Andreas Hertle, Gerhard Lakemeyer, Wolfram Burgard, and Bernhard Nebel, 'Towards Deliberative Active Perception using Persistent Memory', in *Proc. IROS 2013 Workshop on AI-based Robotics*, (2013).
[12] David Rajaratnam, Hector J. Levesque, Maurice Pagnucco, and Michael Thielscher, 'Forgetting in action', in *Proc. KR*, (2014).
[13] Raymond Reiter, 'On closed world data bases', in *Logic and Data Bases*, eds., Herve Gallaire and Jack Minker, (1978).
[14] Raymond Reiter, *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*, The MIT Press, 2001.
[15] Richard Scherl and Hector J. Levesque, 'Knowledge, action, and the frame problem', *Artificial Intelligence*, **144**(1–2), (2003).
[16] Steven Shapiro, Maurice Pagnucco, Yves Lespérance, and Hector J. Levesque, 'Iterated belief change in the situation calculus', *Artificial Intelligence*, **175**(1), (2011).
[17] Moritz Tenorth and Michael Beetz, 'KnowRob – A Knowledge Processing Infrastructure for Cognition-enabled Robots. Part 1: The KnowRob System', *Int. Journal of Robotics Research (IJRR)*, (2013).
[18] Michael Thielscher, 'Representing the knowledge of a robot', in *Proc. KR*, (2000).

[8] http://www.hybrid-reasoning.org/projects/c1