To appear in the *Journal of Experimental & Theoretical Artificial Intelligence*
Vol. 00, No. 00, Month 20XX, 1–19

Sensor Fusion in the Epistemic Situation Calculus

Christoph Schwering[a,†] and Tim Niemueller[a] and Gerhard Lakemeyer[a] and Nichola Abdo[b] and Wolfram Burgard[b]

[a] Knowledge-Based Systems Group, RWTH Aachen University, Germany
[b] Autonomous Intelligent Systems, University of Freiburg, Germany

(*Received 00 Month 20XX; final version received 00 Month 20XX*)

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.

Keywords: actions and knowledge; epistemic situation calculus; sensor fusion

1. Introduction

Even for supposedly straightforward tasks a robot needs to perform complex perception to gather sufficient knowledge about the environment and the relevant objects. Imagine a table with several items on it, like a coffee pot, a pack of sugar, and a mug, as depicted in Figure 1. Depending on her viewpoint, even a human observer may obtain only a partial or perhaps even mistaken view of this scene. Likewise, a robot equipped with an RGB-D camera, which provides color and distance of each pixel, can only perceive those parts of objects which face the camera. For example, in Figure 1 the sugar pack is not visible to the robot because it is occluded by the coffee pot. Similarly, the robot cannot see the mug's handle because it is at the back side. To obtain complete information about the objects on the table, the robot, just like a human, needs to observe the scene from different viewpoints and as the circumstances require even needs to inspect some objects more closely. This is called *active perception* because the robot needs to physically act – for example, move around the table – to obtain new information (Bajcsy, 1988). Active perception goes beyond passive sensing such as robot localization where new information about the robot's position is made available continuously by the navigation subsystem without active intervention by the high-level control of the robot.

To perform active perception reasonably, a robot needs to reason about what it knows and what it can do to obtain additional knowledge. Incomplete and incorrect sensing results need to be combined consistently and find their way into the robot's knowledge. In this paper, we address the outlined problem in the situation calculus. We present a variant

†Corresponding author. Email: schwering@kbsg.rwth-aachen.de
This work was supported by the German National Science Foundation (DFG) under Grant FOR 1513. The first author was partially supported by a DAAD doctoral scholarship.

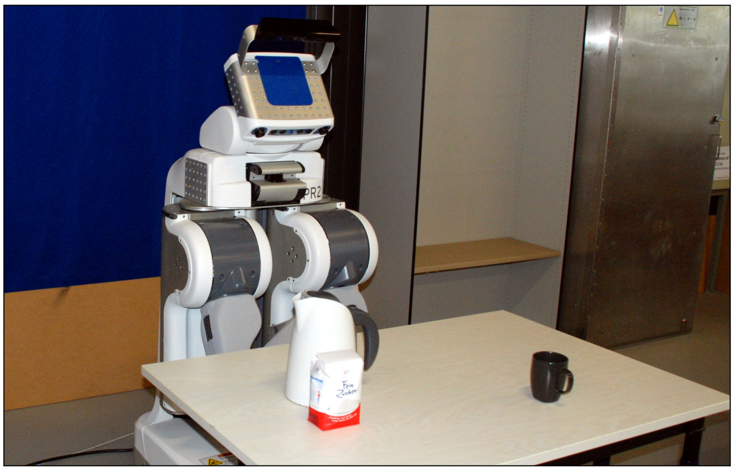

Figure 1. A PR2 robot looking at a table. From the current perspective, the sugar pack is occluded by the coffee pot and the handle of the mug is not visible, causing the robot to confuse it with a cup. From a different perspective, however, the robot would see the sugar pack and recognize the mug correctly.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

of the modal epistemic situation calculus (Lakemeyer & Levesque, 2004), which supports reasoning about incomplete and incorrect sensings. In fact, in our new logic sensing actions have no immediate effect on knowledge to avoid inconsistencies; 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 viewpoints, the robot could fuse these sensings. It then believes[1] that some objects are on the table, but it does not rule out that there may be more. After further sensings and fusion actions from yet more viewpoints, the robot will eventually conclude that it has seen all the objects on the table. In order to reflect this in its knowledge base, the robot will issue an explicit close action, after which it believes that only those objects are on the table which it has seen so far. To undo the epistemic effects of actions, we also incorporate a notion of forgetting. This is useful, for example, after completing a task involving the kitchen table and when it is no longer required to remember which objects were on the table. Likewise, it may be necessary to undo the epistemic effect of a close action in case more objects are found on the table after all.

The rest of the paper is organized as follows. In the next section we present the new logic $\mathcal{ESF}$ and show some of its properties. In Section 3 we model two different scenarios with $\mathcal{ESF}$. While the first example is meant to familiarize the reader 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 based on the modal epistemic situation calculus $\mathcal{ES}$ proposed by Lakemeyer and Levesque (2004). While they later proposed an extended version in (Lakemeyer & Levesque, 2011), we refer to the original logic to simplify the presentation.

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

2.1. *The Language*

The language $\mathcal{ESF}$ consists of formulas over *predicates* and *terms*. To ease the presentation, we only allow *rigid* terms and *fluent* predicates. That is, actions have no effect on the interpretation of a term, but they may affect the truth value of predicates. We assume a countably infinite supply of function and predicate symbols of every arity.

Then 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.

A *ground* term is a term that does not mention variables. We denote the set of all ground terms as R and the set of all finite sequences of ground terms including the empty sequence $\langle\rangle$ as R^*. We often use r and z to denote elements of R and R^*, respectively.

The set of 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 $\mathbf{S}^n_t\alpha$, $\mathbf{K}\alpha$, $\mathbf{O}\alpha$ are formulas.

We read $[t]\alpha$ as "α holds after action t" and $\Box\alpha$ as "α holds after any sequence of actions." $\mathbf{S}^n_t\alpha$ is read as "the nth from last occurrence of action t sensed α," $\mathbf{K}\alpha$ as "α is known," and $\mathbf{O}\alpha$ as "α is all that is known." The purpose of $\mathbf{S}^n_t\alpha$ is to reason about one specific sensing result, whereas $\mathbf{K}\alpha$ refers to knowledge that is the result of fusing previous sensings. $\mathbf{O}\alpha$ is typically used to express that a knowledge base is all the agent knows initially (H. J. Levesque & Lakemeyer, 2001).

We will use $\vee$, $\exists$, $\supset$, $\equiv$, FALSE, TRUE as the usual abbreviations. We let $\mathbf{K}_{\mathbf{if}}\alpha$ stand for $(\mathbf{K}\alpha \vee \mathbf{K}\neg\alpha)$, which is read as "it is known whether or not α."

To ease notation, we consider free variables as being implicitly universally quantified with maximal scope unless noted otherwise. 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. When we omit brackets, the operator precedence in increasing order is: $\Box$, $\forall$, $\exists$, $\equiv$, $\supset$, $\vee$, $\wedge$, $\mathbf{K}_{\mathbf{if}}$, $\mathbf{K}$, $\mathbf{O}$, $\mathbf{S}^n_t$, $[t]$, $\neg$. For example, $\Box[a]F(x) \equiv \gamma_F$ stands for $\forall a.\forall x.\Box(([a]F(x)) \equiv \gamma_F)$. We use sans-serif font for function symbols, like sugar.

There are three special 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 action a.

The precise meaning of these predicates is application-dependent, and therefore is not fixed in the logic but part of the knowledge base (cf. Section 2.4). For example, a background theory might stipulate that $Poss(\mathsf{move})$ is true only if the agent would not bump into a wall when it moves. When a robot looks at a table, $SR(\mathsf{sense}, \mathsf{cup})$ could represent that it perceives a cup. After looking at the table for a while, the robot might come to believe it has seen everything there is to see: for all x, unless it believes $On(x)$, it then believes $\neg On(x)$. This is called a *local closed-world assumption* on $On(x)$. This effect is achieved in $\mathcal{ESF}$ by an action r if $CW(r, x)$ is defined as $On(x)$.

A formula without free variables is a *sentence*. α^x_t denotes the result of substituting t for all free occurrences of x in α. A formula with no modal operators ($[t]$, $\Box$, $\mathbf{S}^n_t$, $\mathbf{K}$, $\mathbf{O}$) and no special predicates ($Poss$ or SR) is called a *fluent* formula. A *fusion* formula has a single free variable a and mentions no $[t]$, $\Box$, $\mathbf{K}$, or $\mathbf{O}$. We denote fusion formulas by ϕ.

2.2. *The Semantics*

In the possible-worlds style semantics of $\mathcal{ESF}$, truth of a sentence is defined wrt

- a world w defining the extension of predicates after any number of actions;
- an epistemic state e representing what is known initially;
- a sensing history h, which memorizes the sensing results of all actions;
- a fusion formula ϕ, which defines how sensing results turn into knowledge;[2]
- a sequence of executed actions z.

We write $\phi, e, w, h, z \models \alpha$ to denote that ϕ, e, w, h, and z satisfy sentence α. Before turning to the semantic rules for the various language constructs, let us briefly discuss our choice of domain of discourse and go over the components of the left-hand side of the satisfaction relation.

As in (Lakemeyer & Levesque, 2004), we choose R as the fixed domain of discourse for all worlds. In essence, this is like using the Herbrand universe as domain of discourse, which also has the effect of making the unique names assumption for all ground terms. Alternatively, we can think of R as a set of *standard names,* which are isomorphic to a fixed, countably infinite domain. As we will see, this choice greatly simplifies the semantics, as we can interpret quantifiers substitutionally and leads to a simple treatment of quantifying-in. While we acknowledge that there are philosophical arguments against substitutional quantification (see, for example, (Kripke, 1976)), we believe its simplicity outweighs its possible disadvantages for our purposes.

The parameter $z \in R^*$ is the sequence of actions executed so far. Intuitively, we begin with the empty sequence $\langle\rangle$ and then deterministically advance with each executed action r to $z \cdot r$. The action sequence z can be considered the world history since the initial situation $\langle\rangle$.[3] A world w maps each ground atomic sentence $P(r_1, \ldots, r_k)$ (for $r_i \in R$) and ground sequence of actions $z \in R^*$ to a truth value $w[P(r_1, \ldots, r_k), z] \in \{0, 1\}$. An epistemic state e is a set of possible worlds, which determine the agent's knowledge. A sensing history h maps each action sequence $z \cdot r$ (for $z \in R^*$ and $r \in R$) to a set of worlds $h(z, r)$, which intuitively contains all those worlds which agree with the sensing result of r after z. Lastly, the fusion formula ϕ^a_r holds iff the given world satisfies the sensor fusion performed by action $r \in R$.

When writing $\phi, e, w, h, z \models \alpha$, we often omit z and h when $z = \langle\rangle$. We allow ourselves to omit further parameters when they are irrelevant to the truth of α. For example, when α is a fluent sentence, we may omit ϕ, e, and h.

2.2.1. Semantics of the Classical Part of the Language

Defining the semantics of the classical, non-modal part of the language is straightforward and works as expected. In particular, note the interpretation of quantifiers by substituting ground terms for the quantified variable.

(1) $\phi, e, w, h, z \models P(r_1, \ldots, r_m)$ iff $w[P(r_1, \ldots, r_m), z] = 1$ for ground terms r_i;
(2) $\phi, e, w, h, z \models (r = r')$ iff r and r' are identical ground terms;
(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 $\phi, e, w, h, z \models \alpha^x_r$ for all $r \in R$.

[2]The fusion formula ϕ is part of the model for technical reasons. Unlike $\mathcal{ES}$ (Lakemeyer & Levesque, 2011) and the Scherl–Levesque (Scherl & Levesque, 2003) 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.
[3]z is the semantic counterpart of the situation terms in the classical situation calculus (Reiter, 2001).

Recall that the connectives $\vee$, $\exists$, $\supset$, and $\equiv$ are just abbreviations.

2.2.2. Actions and Sensing

We now turn to the semantics of actions and their sensing results. Each action r may in principle yield countably infinitely many sensing results, namely those terms r' for which $SR(r, r')$ comes out true. For example, each of these sensing results may represent an object seen on the table. Semantically, this corresponds to the set of all worlds w' which agree on the sensing result purported by the real world of w, that is, $w'[SR(r, r'), z] = w[SR(r, r'), z]$ where z is the current situation. These sets are memorized in the sensing history h. We write $h_r^{w,z}$ for the sensing history h updated by the sensing results of r after z, which is defined as follows:

$$h_r^{w,z}(z, r) \doteq \{w' \mid w'[SR(r, r''), z] = w[SR(r, r''), z] \text{ for all } r'' \in R\};$$
$$h_r^{w,z}(z', r') \doteq h(z', r') \text{ for all } z' \cdot r' \neq z \cdot r.$$

We use $h_{z'}^{w,z}$ as a shorthand for h updated with the sensing results throughout $z' \in R^*$. The semantics of action execution can then be defined in the following way:

(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 $z' \in R^*$.

The memorized sensing results can be accessed through the $\mathbf{S}_t^n$ modality. 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, $\langle \mathsf{f}, \mathsf{g}, \mathsf{f}, \mathsf{h}, \mathsf{f} \rangle|_{\mathsf{f}}^2 = \langle \mathsf{f}, \mathsf{g} \rangle$. Then we obtain the following rule:

(8) $\phi, e, w, h, z \models \mathbf{S}_r^n\alpha$ iff $|z|_r \geq n$ and $\phi, e, w', h, z \models \alpha$ for all $w' \in e \cap h(z|_r^n, r)$.

Observe that the sensing history is intersected with e and α is evaluated wrt z. As we will see in Section 3.1, this ensures that sensing results are adequately projected into the current situation. Intuitively, this is because e represents the agent's knowledge about the domain's dynamics. Also notice that if there is no nth last occurrence of r yet, $\mathbf{S}_r^n\alpha$ is vacuously false.

2.2.3. Knowledge

So far, we have not established any connection between knowledge and sensing. Instead, sensing results are just memorized one by one in the sensing history. We now define how knowledge is produced from sensings. The idea is that some later action may fuse (possibly contradicting) earlier sensing results.

Semantically, producing new knowledge means that fewer possible worlds need to be considered, that is, the epistemic state e needs to be filtered. To characterize this filtering, we define an operator $e{\downarrow}^{\phi,h,z}$ that retains only those worlds from e that agree with the fusion specified by ϕ of the sensing results h throughout z:[4]

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

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

Condition (2) requires that w satisfies the fusion formula ϕ_r^a. Recall that a fusion formula has a single free variable a, which represents the latest action, and usually a few subformulas involving $\mathbf{S}_t^n$ occur as well. For example, suppose a fluent predicate $D(x)$ expresses that a robot's distance to a wall currently is x. The robot senses the distance through the action sense and the action fuse shall fuse the latest two sensings by taking their average. The corresponding fusion formula is $a = \mathsf{fuse} \supset \exists x_1.\exists x_2.\mathbf{S}^1_{\mathsf{sense}} D(x_1) \wedge \mathbf{S}^2_{\mathsf{sense}} D(x_2) \wedge D(\frac{x_1+x_2}{2})$. Then only those worlds are kept in $e{\downarrow}^{\phi,h,z\cdot\mathsf{fuse}}$ where the robot's distance is the average of the latest sensed distances.

Condition (3) makes a closed-world assumption for $CW(r,x)$ for given r and variable x. For example, suppose a robot looks at a table for objects and $On(x)$ holds iff object x is on the table. After inspecting the table from several perspectives, the robot may want to assume it has seen every object on it. This is reflected by an action close which has no physical effect but makes a closed-world assumption for $On(x)$ through the axiom $CW(\mathsf{close},x) \equiv On(x)$. We examine this example in Section 3.2.

With that definition in hand, the subjective semantics is rather straightforward:

(9) $\phi, e, w, h, z \models \mathbf{K}\alpha$ iff $\phi, e, w', h, z \models \alpha$ for all $w' \in e{\downarrow}^{\phi,h,z}$;

(10) $\phi, e, w, h, z \models \mathbf{O}\alpha$ iff $\phi, e, w', h, z \models \alpha$ iff $w' \in e{\downarrow}^{\phi,h,z}$.

The only difference between $\mathbf{K}$ and $\mathbf{O}$ the "only-if" direction, which intuitively maximizes the epistemic state e. That way, only-knowing can be used to concisely specify the agent's knowledge and implicitly also what she does not know.

This completes the semantics of $\mathcal{ESF}$ for now; we defer the definition of *forgetting* to Section 2.6.

We remark that the closed-world assumption as defined in $e{\downarrow}^{\phi,h,z}$ is about *complete* knowledge. Suppose the agent believes $P(r) \vee P(r')$ in situation z, that is, $e{\downarrow}^{\phi,h,z}$ contains at least one world which satisfies only $P(r)$ and another which satisfies only $P(r')$. When we make a closed-world assumption for $P(x)$ (through some action close with $CW(\mathsf{close},x) \equiv P(x)$), then there is no world left in $e{\downarrow}^{\phi,h,z\cdot\mathsf{close}}$ and the agent's knowledge is hence inconsistent. This problem could be addressed through the *generalized* closed-world assumption (Minker, 1982).

We conclude this subsection with the definition of entailment: A set of sentences Σ entails α wrt a fusion formula ϕ (written $\Sigma \models_\phi \alpha$) iff for all e, w, and h, if $\phi, e, w, h, \langle\rangle \models \sigma$ for all $\sigma \in \Sigma$, then $\phi, e, w, h, \langle\rangle \models \alpha$. 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 $\{\} \models \alpha$.

2.3. *Some Properties*

In this subsection, we investigate introspection and the local closed-world assumption in $\mathcal{ESF}$. To begin with, $\mathcal{ESF}$ is fully introspective:

Theorem 2.1: $\models \Box\mathbf{K}\alpha \supset \mathbf{K}\mathbf{K}\alpha$ *and* $\models \Box\neg\mathbf{K}\alpha \supset \mathbf{K}\neg\mathbf{K}\alpha$.

Proof. For positive introspection, suppose $\phi, e, w, h, z \models \mathbf{K}\alpha$. By rule 9, $\phi, e, w', h, z \models \alpha$ for all $w' \in e{\downarrow}^{\phi,h,z}$. Then, obviously, $\phi, e, w'', h, z \models \alpha$ for all $w' \in e{\downarrow}^{\phi,h,z}$ and for all $w'' \in e{\downarrow}^{\phi,h,z}$. Applying rule 9 twice yields $\phi, e, w, h, z \models \mathbf{K}\mathbf{K}\alpha$.

For negative introspection, suppose $\phi, e, w, h, z \models \neg\mathbf{K}\alpha$. By rule 9, $\phi, e, w', h, z \models \neg\alpha$ for some $w' \in e{\downarrow}^{\phi,h,z}$. Then, obviously, $\phi, e, w'', h, z \models \neg\alpha$ for some $w'' \in e{\downarrow}^{\phi,h,z}$ for all $w' \in e{\downarrow}^{\phi,h,z}$. Applying rule 9 twice yields $\phi, e, w, h, z \models \mathbf{K}\neg\mathbf{K}\alpha$. □

The following theorem expresses that a given action a performs a local closed-world assumption for $CW(a,x)$, that is, afterwards either $CW(a,x)$ is known or $\neg CW(a,x)$ is known for all x. In the above example this means that after action close, we know for each object r' whether or not $On(r')$ is true, that is, if r' is on the table or not. Proviso for the theorem is that the performed action does not affect the truth value of CW:

Theorem 2.2: $\models \Box\mathbf{K}([a]CW(a,x) \equiv CW(a,x)) \supset [a]\mathbf{K}_{\mathbf{if}}CW(a,x)$.

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

If $\phi, e, w, h^w_z, z \models \mathbf{K}CW(r,r')$, condition (3) 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 the truth of $CW(r,r')$, it follows $\phi, e, w, h^w_{z\cdot r}, z\cdot r \models \mathbf{K}CW(r,r')$.

Now suppose $\phi, e, w, h^w_z, z \models \neg\mathbf{K}CW(r,r')$. Then condition (3) requires $w'[CW(r,r'), z] = 0$ for all $w' \in e\!\downarrow^{\phi, h^w_{z\cdot r}, z\cdot r}$. As r does not change the truth of $CW(r,r')$, it follows $\phi, e, w, h^w_{z\cdot r}, z\cdot r \models \mathbf{K}\neg CW(r,r')$. □

2.4. *Basic Action Theories*

We are mostly interested in evaluating queries after a sequence of actions with respect to a knowledge base, the so called *projection problem*. Formally, this problem can be expressed by[5]

$$\Sigma \wedge \mathbf{O}\Sigma' \models_\phi [r_1]\ldots[r_k]\mathbf{K}\alpha,$$

where Σ and Σ' are knowledge bases that represent what is *actually* true (Σ) or what is *believed* to be true (Σ'), ϕ is a fusion formula, $r_1, \ldots, r_k$ are actions, and α is a query which may or may not be believed. The knowledge bases we consider here are a variant of Reiter's basic action theories (Reiter, 2001). A basic action theory not only determines what is true or believed to be true initially, but also stipulates how these beliefs change when actions are being executed. More precisely, a basic action theory 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 theory Σ and the theory 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:

- 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 $\Box SR(a,x) \supset \alpha$;
- Σ'_{sense} is a singleton sentence of the form $\Box SR(a,x) \equiv \alpha$;
- Σ'_{close} is a singleton sentence of the form $\Box CW(a,x) \equiv \alpha$;

where all α's are fluent formulas. The sentences in Σ_{post} are called *successor state axioms*. Successor state axioms are of the form

$$\Box[a]F(\vec{x}) \equiv \gamma^+_F(a,\vec{x}) \vee F(\vec{x}) \wedge \neg\gamma^-_F(a,\vec{x})$$

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

where $\gamma_F^{\pm}(a,\vec{x})$ capture the positive and negative effects, respectively, of a on $F(\vec{x})$. This pattern is key to Reiter's solution to the frame problem (Reiter, 2001). Observe that Σ and Σ' not only may 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}, which is just an implication, merely constrains the possible sensing results to rule out implausible values. An actual sensing thus may yield any set of results that satisfies that constraint.

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

For example basic action theories and queries we refer to Section 3.

2.5. *Relationship to the Logic $\mathcal{ES}$*

We now relate $\mathcal{ESF}$ to its ancestor $\mathcal{ES}$ (Lakemeyer & Levesque, 2004). In $\mathcal{ES}$, actions may yield a binary sensing result, and the agent immediately knows that outcome. In this section we show that the projection problem in $\mathcal{ES}$ is reducible to the projection problem $\mathcal{ESF}$.

The language $\mathcal{ES}$ is similar to $\mathcal{ESF}$, except that $\mathbf{K}$ and $\mathbf{O}$ are termed *Know* and *OKnow*, respectively, and there is no counterpart for $\mathbf{S}_t^n$. Instead of our *SR* predicate, $\mathcal{ES}$ features a special unary predicate *SF* to represent an action's binary sensing result. The semantics of the objective part of the language matches rules 1–7 with ϕ and h being no longer required. To characterize what is known after a sequence of actions, Lakemeyer and Levesque use a relation $w \simeq_z w'$ to express that w and w' agree on the sensing results throughout z. It is defined inductively on z:

- $w \simeq_{\langle\rangle} w'$ for all w';
- $w \simeq_{z\cdot r} w'$ iff $w \simeq_z w'$ and $w[SF(r), z] = w'[SF(r), z]$.

Then the subjective semantics in $\mathcal{ES}$ is as follows:

- $e, w, z \models_{\mathcal{ES}} Know(\alpha)$ iff for all $w' \simeq_z w$, if $w' \in e$ then $e, w', z \models_{\mathcal{ES}} \alpha$;
- $e, w, z \models_{\mathcal{ES}} OKnow(\alpha)$ iff for all $w' \simeq_z w$, $w' \in e$ iff $e, w', z \models_{\mathcal{ES}} \alpha$.

We translate $\mathcal{ES}$ formulas α to $\mathcal{ESF}$. We map any formula α of $\mathcal{ES}$ into a formula α^* of $\mathcal{ESF}$, where α^* is α with all occurrences of *Know* and *OKnow* replaced by $\mathbf{K}$ and $\mathbf{O}$, respectively. We further align the sensing models of $\mathcal{ES}$ and $\mathcal{ESF}$ as follows. First, we assert in the knowledge base $\forall a. SR(a, \mathsf{true}) \equiv SF(a)$ where $\mathsf{true} \in R$, that is, $\mathcal{ESF}$'s *SR* predicate is restricted to represent binary sensing result. Second, we use as fusion formula $\mathbf{S}_a^0 SR(a, \mathsf{true}) \equiv SR(a, \mathsf{true})$, so that the sensing history will contain just those worlds that agree with the real world's *SF* truth value. Finally we also assert $\forall a. CW(a) \equiv \text{TRUE}$, so that no action performs any closure. Then the following theorem says how the projection problem in $\mathcal{ES}$ can be reduced to the projection problem in $\mathcal{ESF}$:

Theorem 2.3: *Let Γ and Γ' be $\mathcal{ES}$ knowledge bases without Know and OKnow. Let α be an $\mathcal{ES}$ formula without OKnow and that does not mention* $[\mathsf{forget}(a)]$. *Let*

$$\phi \doteq \mathbf{S}_a^0 SR(a, \mathsf{true}) \equiv SR(a, \mathsf{true});$$
$$\psi \doteq (\forall a. SR(a, \mathsf{true}) \equiv SF(a)) \wedge (\forall a. CW(a) \equiv \text{TRUE}).$$

Then $OKnow(\Gamma' \wedge \psi) \wedge \Gamma \wedge \psi \models_{\mathcal{ES}} \alpha$ iff $\mathbf{O}(\Gamma'^* \wedge \psi) \wedge \Gamma^* \wedge \psi \models_\phi \alpha^*$.

Proof. Due to the equivalent semantic rules 1–7, for any formula β without *Know*, *OKnow*, $\models_{\mathcal{ES}} \beta$ iff $\models \beta^*$. Therefore, $w, \langle\rangle \models_{\mathcal{ES}} \Gamma \wedge \psi$ iff $w, \langle\rangle \models_\phi \Gamma'^* \wedge \psi$, and by the semantics of $OKnow$ and $\mathbf{O}$, also $e, w, \langle\rangle \models OKnow(\Gamma' \wedge \psi) \wedge \Gamma \wedge \psi$ iff $\phi, e, w, h, \langle\rangle \models \mathbf{O}(\Gamma'^{*} \wedge \psi) \wedge \Gamma^{*} \wedge \psi$.

Let $e, w, \langle\rangle \models_{\mathcal{ES}} OKnow(\Gamma' \wedge \psi) \wedge \Gamma \wedge \psi$. We need to show that $e, w, \langle\rangle \models_{\mathcal{ES}} \alpha$ iff $\phi, e, w, h, \langle\rangle \models \alpha^{*}$.

We need one more observation: $w' \in e$ and $w' \simeq_{z} w$ iff $w' \in e{\downarrow}^{\phi,h,z}$ where h was built up through z (**). We prove this by induction on z. The base case $z = \langle\rangle$ follows from the definitions. Consider $z \cdot r$ for the induction step. Then $w' \in e$ and $w' \simeq_{z \cdot r} w$ iff $w' \in e$ and $w'[SF(r), z] = w[SF(r), z]$ and $w' \simeq_{z} w$ iff (by induction hypothesis) $w'[SF(r), z] = w[SF(r), z]$ and $w' \in e{\downarrow}^{\phi,h,z} \subseteq e$ iff (by axiom ψ) $w'[SR(r, \mathsf{true}), z] = w[SR(r, \mathsf{true}), z]$ and $w' \in e{\downarrow}^{\phi,h,z}$ iff (since h was built up through z) $w'[SR(r, \mathsf{true}), z] = w''[SR(r, \mathsf{true}), z]$ for all $w'' \in e \cap h(z, r)$ and $w' \in e{\downarrow}^{\phi,h,z}$ iff (by axiom ϕ) $e, h, w', z \models \phi_{r}^{a}$ and $w' \in e{\downarrow}^{\phi,h,z}$ iff $w' \in e{\downarrow}^{\phi,h,(z \cdot r)}$.

We now prove $e, w, z \models_{\mathcal{ES}} \alpha$ iff $\phi, e, w, h, z \models \alpha^{*}$ where h was built up through z. For the first base case, consider a ground atom $P(\vec{r})$. Then $e, w, z \models_{\mathcal{ES}} F(\vec{r})$ iff $w[F(\vec{r}), z] = 1$ iff $\phi, e, w, h, z \models (F(\vec{r}))^{*}$. We skip the base case ($r = r'$) and the induction steps for $\wedge$, $\neg$, and $\forall$ as they follow immediately from the common semantics. For $[r]\alpha$, we have $e, w, z \models_{\mathcal{ES}} [r]\alpha$ iff $e, w, z \cdot r \models_{\mathcal{ES}} \alpha$ iff (by induction hypothesis) $\phi, e, w, h_{r}^{w,z}, z \cdot r \models \alpha^{*}$ iff $\phi, e, w, h, z \models ([r]\alpha)^{*}$. The induction step for $\Box\alpha$ is analogous. For $Know(\alpha)$, we have $e, w, z \models_{\mathcal{ES}} Know(\alpha)$ iff for all $w' \in e$ with $w' \simeq_{z} w$, $e, w', z \models_{\mathcal{ES}} \alpha$ iff (by induction hypothesis) for all $w' \in e$ with $w' \simeq_{z} w$, $\phi, e, w', h, z \models \alpha^{*}$ iff (by (**)) for all $w' \in e{\downarrow}^{\phi,h,z}$, $\phi, e, w', h, z \models \alpha^{*}$ iff $\phi, e, w, h, z \models (\mathbf{K}\alpha)^{*}$. $\Box$

We remark that a more general version of the theorem, such as $\models_{\mathcal{ES}} \alpha$ iff $\models \alpha^{*}$, does not hold, due to the different sensing models of $\mathcal{ES}$ and $\mathcal{ESF}$. For example, let $\Gamma = \{[\mathsf{sense}]On(x) \equiv On(x), SF(\mathsf{sense}) \equiv On(\mathsf{cup})\}$. In $\mathcal{ES}$, this means that the action sense does not affect the truth of $On(x)$, but it produces knowledge if $On(\mathsf{cup})$ is true. Therefore, $\models_{\mathcal{ES}} OKnow(\Gamma) \wedge \Gamma \wedge On(\mathsf{cup}) \supset [\mathsf{sense}]Know(On(\mathsf{cup}))$: Suppose $e, w \models_{\mathcal{ES}} OKnow(\Gamma) \wedge \Gamma \wedge On(\mathsf{cup})$. Initially the worlds in e do not agree on the truth value of $On(\mathsf{cup})$. But after sense, only the worlds $w' \in e$ with $w' \simeq_{\mathsf{sense}} w$ are considered, that is, $w'[On(\mathsf{cup}), \langle\rangle] = w[On(\mathsf{cup}), \langle\rangle] = 1$, and therefore $On(\mathsf{cup})$ is known. In $\mathcal{ESF}$, this does not hold for several reasons. Firstly, sensing is not turned into knowledge immediately in $\mathcal{ESF}$, so an appropriate fusion formula is necessary. Secondly, the epistemic effects of action sense depends on $SR(\mathsf{sense}, x)$ and $CW(\mathsf{sense}, x)$, which need to be axiomatized in Γ to achieve reasonable results. Hence, if $\phi, e, w, h, z \models \mathbf{O}\Gamma \wedge \Gamma \wedge On(\mathsf{cup})$, then $e{\downarrow}^{\phi,h,z}$ may be any subset of e, and therefore no knowledge is gained after sense.

2.6. *Forgetting*

In many settings it may be desirable to forget some information, for example, to revoke a sensor fusion or a local closed-world assumption. Rajaratnam, Levesque, Pagnucco, and Thielscher (2014) proposed an extension of the Scherl and Levesque's (2003) approach to knowledge in the situation calculus. We adopt their ideas for our logic $\mathcal{ESF}$ in this subsection.

We introduce a special action term $\mathsf{forget}(r)$ where r denotes another action term. The idea is that $\mathsf{forget}(r)$ undoes the epistemic effect of the last occurrence of action r. To this end, we replace condition (1) of $e{\downarrow}^{\phi,h,z \cdot r}$ with

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

The trick is simply to skip r in the filtering of $e{\downarrow}^{\phi,w,z \cdot \mathsf{forget}(r)}$. Thus the worlds rendered impossible by r survive the filtering step and the epistemic effect of r is thus undone.

The following theorem shows that the effect of an action a is undone by $\mathsf{forget}(a)$ with the proviso that a and $\mathsf{forget}(a)$ have no physical effect and $\mathsf{forget}(a)$ has no other epistemic effect:

Theorem 2.4: *Let α be a sentence and let ϕ be a fusion formula. Let β abbreviate the formula $([a][\mathsf{forget}(a)]\alpha \equiv \alpha) \land [a]\phi^{a}_{\mathsf{forget}(a)} \land [a]\forall x.CW(\mathsf{forget}(a), x)$ where a is a free variable. Then the epistemic effect of action a can be revoked by $\mathsf{forget}(a)$:*

$$\models_\phi \Box\mathbf{K}\beta \land \neg\mathbf{K}\alpha \land [a]\mathbf{K}\alpha \supset [a][\mathsf{forget}(a)]\neg\mathbf{K}\alpha$$

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 (1') skips action conditions (2) and (3) for r, and conditions (2) and (3) 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 (Rajaratnam et al., 2014). It is straightforward to generalize to more actions between r and $\mathsf{forget}(r)$.

We remark that our definition and proof is much simpler than their counterparts in (Rajaratnam et al., 2014). This is because Rajaratnam et al. need to carefully manage several accessibility relations after each action to keep track of the reachable situations at different points in time.

3. Examples

In this section we model two scenarios as basic action theories and show a few properties. The first example is meant to familiarize with $\mathcal{ESF}$. We model a variant of the running example from (Lakemeyer & Levesque, 2011) and show how sensings can be fused by taking their disjunction, which is not possible in Scherl–Levesque style approaches (Lakemeyer & Levesque, 2011; Scherl & Levesque, 2003). In the second example we return to our motivating scenario of a robot investigating objects on a table. We show how to deal with an unknown number of objects and exemplify use of the local closed-world assumption and of forgetting.

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:

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

The robot may move one unit towards the wall (through action move). The appropriate precondition axiom and successor state axiom are:

$$\Sigma_{pre} \doteq \{\Box Poss(a) \equiv (a = \mathsf{move} \supset \neg D(0))\}$$
$$\Sigma_{post} \doteq \{\Box D(x) \equiv \exists x'.a = \mathsf{move} \land D(x') \land x = x' - 1 \lor D(x) \land a \neq \mathsf{move}\}$$

The robot is equipped with a sonar sensor (action sense), which yields intervals of possible distances. In Σ_{sense} we express the constraint that each sensing result is indeed an interval $[x_1, x_2]$. Σ'_{sense} captures that a possible world agrees on a sensing result $[x_1, x_2]$ iff the distance x in that world is in $[x_1, x_2]$. In our basic action theory we hence have the following definitions:

$$\Sigma_{sense} \doteq \{\Box SR(a, y) \supset (a = \mathsf{sense} \supset \exists x_1.\exists x_2.y = [x_1, x_2])\}$$
$$\Sigma'_{sense} \doteq \{\Box SR(a, y) \equiv (a = \mathsf{sense} \supset \exists x_1.\exists x_2.y = [x_1, x_2] \wedge \exists x.D(x) \wedge x_1 \leq x \leq x_2)\}$$

Notice that a sensing result $[x_1, x_2]$ represents disjunctive information. Disjunctive information must be encoded within such a single sensing result, as opposed to a set of sensing results like $SR(\mathsf{sense}, x)$ for all $x_1 \leq x \leq x_2$. 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. In our example, constraining $SR(\mathsf{sense}, x)$ for all $x_1 \leq x \leq x_2$ would hence require each possible world to have *all* distances in $[x_1, x_2]$, which is unreasonable and in fact inconsistent with Σ'_0, which asserts that there is precisely one distance in each possible world.

Since our robot mistrusts its own sensor, we take the fusion (action fuse) 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

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

Notice that $\neg \mathbf{S}^1_{\mathsf{sense}} \neg D(x)$ expresses that x was not ruled out by the last sensing of sense, as it means that there is at least one possible world which satisfies $D(x)$.

We do not use the local closed-world assumption here, so we define:

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

Now we can reason about what is entailed by this basic action theory. To this end, let e, w be such that

$$\phi, e, w \models \Sigma \wedge \mathbf{O}\Sigma' \wedge SR(\mathsf{sense}, [4, 7]) \wedge [\mathsf{sense}]SR(\mathsf{sense}, [3, 6]),$$

that is, the first sense reports the interval $[4, 7]$ and a subsequent sense reports $[3, 6]$. We show the following properties:

The robot believes the distance is in $[3, 7]$ *after sensing twice and fusing both sensings:*

$$\phi, e, w, h \models [\mathsf{sense}][\mathsf{sense}][\mathsf{fuse}]\mathbf{K}(D(x) \equiv 3 \leq x \leq 7)$$

For the proof let z stand for $\langle \mathsf{sense}, \mathsf{sense}, \mathsf{fuse} \rangle$ and let $w' \in e{\downarrow}^{\phi, h^w_z, z}$. Then $e, w', h^w_z, \langle \mathsf{sense}, \mathsf{sense} \rangle \models \phi^a_{\mathsf{fuse}}$ due to condition (2) for $e{\downarrow}^{\phi, h^w_z, z}$, that is, $e, w', h^w_z, \langle \mathsf{sense}, \mathsf{sense} \rangle \models \exists x.D(x) \wedge (\neg \mathbf{S}^1_{\mathsf{sense}} \neg D(x) \vee \neg \mathbf{S}^2_{\mathsf{sense}} \neg D(x))$. We have $e, h^w_z, \langle \mathsf{sense}, \mathsf{sense} \rangle \models \neg \mathbf{S}^2_{\mathsf{sense}} \neg D(x)$ iff $w'[D(x), \langle \mathsf{sense}, \mathsf{sense} \rangle] = 1$ for some $w' \in e \cap \{w' \mid w'[SR(\mathsf{sense}, r), \langle\rangle] = w[SR(\mathsf{sense}, r), \langle\rangle]$ for all $r \in R\}$ iff $4 \leq x \leq 7$. Analogously $e, h^w_z, \langle \mathsf{sense}, \mathsf{sense} \rangle \models \neg \mathbf{S}^1_{\mathsf{sense}} \neg D(x)$ iff $3 \leq x \leq 6$. Thus the property follows.

When the robot moves before fusion, its effect is projected onto the sensing results:

$$\phi, e, w, h \models [\mathsf{sense}][\mathsf{sense}][\mathsf{move}][\mathsf{fuse}]\mathbf{K}(D(x) \equiv 2 \leq x \leq 6)$$

The proof is analogous to the previous one except that $\langle \mathsf{sense}, \mathsf{sense} \rangle$ is replaced with $\langle \mathsf{sense}, \mathsf{sense}, \mathsf{move} \rangle$, which leads to intervals $[2, 5]$ and $[3, 6]$ instead of $[3, 6]$ and $[4, 7]$.

3.2. *Tabletop Object Search*

In our previous work (Niemueller et al., 2013), 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 $On(x)$ to express that object x is on the table and $Is(x, y)$ to say that x is of type y. 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 move). The robot may look on the table (action sense) 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 closed-world assumption (action close).

We proceed to define the objective and the subjective basic action theories 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:

$$\Sigma_0 \doteq \Sigma'_0 \doteq \{L,\ \exists y.Is(x, y),\ Is(x, y) \land Is(x, y') \supset y = y'\}$$

There are no specific preconditions in this scenario:

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

The predicates Is and On shall be rigid, that is, the types of objects remain the same and they remain on or off the table (for the lack of a pickup action in our example). Only the robot's position, represented by the truth of the predicate L, can change. We hence obtain the following successor state axioms:

$$\begin{aligned}\Sigma_{post} \doteq \{&\Box[a]L \equiv (a = \mathsf{move} \land \neg L) \lor (L \land a \neq \mathsf{move}),\\ &\Box[a]Is(x, y) \equiv Is(x, y),\\ &\Box[a]On(x) \equiv On(x)\}\end{aligned}$$

Now we turn to Σ_{sense} and Σ'_{sense}. We know that the sensor only reports objects and

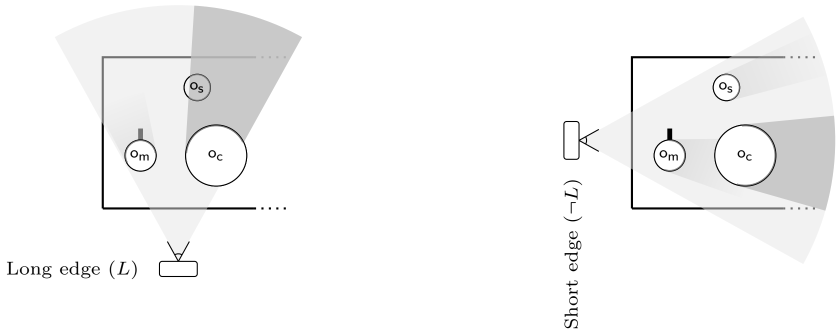

Figure 2. Tabletop with a mug $\mathsf{o_m}$, a sugar box $\mathsf{o_s}$, and a coffee pot $\mathsf{o_c}$ from two different perspectives. Light gray cones denote horizontal viewing angle, dark gray regions represent sensor shadows. From the first perspective, $\mathsf{o_s}$ is partly occluded by $\mathsf{o_c}$. $\mathsf{o_m}$'s handle is only visible from the second perspective.

their types, so we constrain the reported sensor values accordingly:

$$\Sigma_{sense} \doteq \{\Box SR(a,x) \supset (a = \mathsf{sense} \supset \exists x'.x = \mathsf{obj}(x') \lor \\ \exists x'.\exists y.x = \mathsf{type}(x',y))\}$$

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

$$\Sigma'_{sense} \doteq \{\Box SR(a,x) \equiv (a = \mathsf{sense} \supset \exists x'.x = \mathsf{obj}(x') \land On(x') \lor \\ \exists x'.\exists y.x = \mathsf{type}(x',y) \land Is(x',y))\}$$

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.
- If an object was recognized as a y in either sensing, we believe it is a y, modulo one constraint: if y 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 overrule sensing results cup.

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

$$\alpha \doteq \forall x.\mathbf{S}^1_{\mathsf{sense}} On(x) \lor \mathbf{S}^2_{\mathsf{sense}} On(x) \supset On(x)$$

$$\beta \doteq \forall x.\forall y.(\mathbf{S}^1_{\mathsf{sense}} Is(x,y) \land (x = \mathsf{cup} \supset \neg\mathbf{S}^2_{\mathsf{sense}} Is(x,\mathsf{mug}))) \lor \\ (\mathbf{S}^2_{\mathsf{sense}} Is(x,y) \land (x = \mathsf{cup} \supset \neg\mathbf{S}^1_{\mathsf{sense}} Is(x,\mathsf{mug}))) \supset Is(x,y)$$

Then we define our fusion formula as

$$\phi \doteq a = \mathsf{fuse} \supset \alpha \land \beta$$

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

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

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.

For the remainder of this subsection we let

$$\phi, e, w, h \models \Sigma \wedge \mathbf{O}\Sigma'.$$

Imagine that the real world w is as depicted in Figure 2 with three objects $\mathsf{o_m}, \mathsf{o_c}, \mathsf{o_s}$ on the table, where $\mathsf{o_m}$ is a mug (with a handle), $\mathsf{o_c}$ is a large coffee pot, and $\mathsf{o_s}$ is a sugar box:

$$w \models On(\mathsf{o_m}) \wedge On(\mathsf{o_c}) \wedge On(\mathsf{o_s}) \wedge Is(\mathsf{o_m}, \mathsf{mug}) \wedge Is(\mathsf{o_c}, \mathsf{coffee}) \wedge Is(\mathsf{o_s}, \mathsf{sugar})$$

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. In logic:

$$\begin{aligned} w \models \Box SR(\mathsf{sense}, x) \equiv\ & x = \mathsf{obj}(\mathsf{o_m}) \vee x = \mathsf{obj}(\mathsf{o_c}) \vee x = \mathsf{type}(\mathsf{o_c}, \mathsf{coffee}) \vee \\ & (L \wedge x = \mathsf{type}(\mathsf{o_m}, \mathsf{cup})) \vee \\ & (\neg L \wedge (x = \mathsf{obj}(\mathsf{o_s}) \vee x = \mathsf{type}(\mathsf{o_m}, \mathsf{mug}) \vee x = \mathsf{type}(\mathsf{o_s}, \mathsf{sugar}))) \end{aligned}$$

We conclude the example by showing the following properties:

After sensing only from the long side, the robot erroneously thinks $\mathsf{o_m}$ *is a* cup*:*

$$\phi, e, w, h \models [\mathsf{sense}][\mathsf{fuse}]\mathbf{K}Is(\mathsf{o_m}, \mathsf{cup})$$

For the proof let z stand for $\langle\mathsf{sense}\rangle$. We have $e, h_z^w, z \models \mathbf{S}^1_{\mathsf{sense}} Is(\mathsf{o_m}, \mathsf{cup})$ and $e, h_z^w, z \not\models \mathbf{S}^2_{\mathsf{sense}} Is(\mathsf{o_m}, \mathsf{mug})$ since there is just one sensing so far. This and the definition of ϕ^a_{fuse} give that we have $w'[Is(\mathsf{o_m}, \mathsf{cup}), z \cdot \mathsf{fuse}] = 1$ for all $w' \in e{\downarrow}^{\phi, h^w_{z\cdot\mathsf{fuse}}, z\cdot\mathsf{fuse}}$.

After sensing from both sides and forgetting the first fusion, the robot correctly believes $\mathsf{o_m}$ *is a* mug*:*

$$\phi, e, w, h \models [\mathsf{sense}][\mathsf{fuse}][\mathsf{move}][\mathsf{sense}][\mathsf{forget}(\mathsf{fuse})][\mathsf{fuse}]\mathbf{K}Is(\mathsf{o_m}, \mathsf{mug})$$

Let z abbreviate $\langle\mathsf{sense}, \mathsf{fuse}, \mathsf{move}, \mathsf{sense}, \mathsf{forget}(\mathsf{fuse}), \mathsf{fuse}\rangle$ and let z' stand for $\langle\mathsf{sense}, \mathsf{move}, \mathsf{sense}, \mathsf{fuse}\rangle$, which is just z without $\mathsf{forget}(\mathsf{fuse})$ and the corresponding occurrence of fuse. Observe that $e{\downarrow}^{\phi, h^w_z, z} = e{\downarrow}^{\phi, h^w_{z'}, z'}$ because condition (1') means that $\mathsf{forget}(\mathsf{fuse})$ undoes the epistemic effects of the first fuse. Furthermore we have

$e, h_z^w, \langle \mathsf{sense}, \mathsf{move}, \mathsf{sense} \rangle \models \mathbf{S}^1_{\mathsf{sense}} Is(\mathsf{o_m}, \mathsf{mug})$. Thus $w'[Is(\mathsf{o_m}, \mathsf{mug}), z] = 1$ for all $w' \in e{\downarrow}^{\phi, h_z^w, z}$.

After one sensing and fusion, the robot does not believe that $\mathsf{o_m}$ *and* $\mathsf{o_c}$ *are the only objects on the table:*

$$\phi, e, w, h \models [\mathsf{sense}][\mathsf{fuse}]\neg\mathbf{K}(x = \mathsf{o_m} \lor x = \mathsf{o_c} \equiv On(x))$$

There is some $w' \in e$ in which some $r \notin \{\mathsf{o_m}, \mathsf{o_c}\}$ is believed to be on the table as well, that is, $w' \models (x = \mathsf{o_m} \lor x = \mathsf{o_c} \lor x = r \supset On(x)) \land Is(\mathsf{o_m}, \mathsf{cup}) \land Is(\mathsf{o_c}, \mathsf{coffee})$. Since w' satisfies the fusion formula, $w' \in e{\downarrow}^{\phi, h^w_{\langle \mathsf{sense}, \mathsf{fuse} \rangle}, \langle \mathsf{sense}, \mathsf{fuse} \rangle}$. Hence $(x = \mathsf{o_m} \lor x = \mathsf{o_c} \equiv On(x))$ is not known.

After additionally closing the domain, the robot believes there are no objects on the table other than $\mathsf{o_m}$ *and* $\mathsf{o_c}$*:*

$$\phi, e, w, h \models [\mathsf{sense}][\mathsf{fuse}][\mathsf{close}]\mathbf{K}(x = \mathsf{o_m} \lor x = \mathsf{o_c} \equiv On(x))$$

Let z stand for $\langle \mathsf{sense}, \mathsf{fuse}, \mathsf{close} \rangle$ and $w' \in e{\downarrow}^{\phi, h_z^w, z}$. Notice that $w', z \models On(\mathsf{o_m}) \land On(\mathsf{o_c})$ holds because otherwise condition (2) for $e{\downarrow}^{\phi, h_z^w, \langle \mathsf{sense}, \mathsf{fuse} \rangle}$ would be violated. To see that $w'[On(r), z] = 0$ for all $r \notin \{\mathsf{o_m}, \mathsf{o_c}\}$, suppose $w_1, w_2 \in e{\downarrow}^{\phi, h_z^w, \langle \mathsf{sense}, \mathsf{fuse} \rangle}$ and $w_1[On(r), \langle \mathsf{sense}, \mathsf{fuse} \rangle] \neq w_2[On(r), \langle \mathsf{sense}, \mathsf{fuse} \rangle]$. Such w_1 and w_2 exist as argued in the previous property. Then due to condition (3) of $e{\downarrow}^{\phi, h_z^w, z}$, $w'[CW(\mathsf{close}, r), \langle \mathsf{sense}, \mathsf{fuse} \rangle] = 0$, which implies that $w'[On(r), z] = 0$.

After inspecting the table from both sides and fusing these sensings, the robot believes that $\mathsf{o_m}, \mathsf{o_c}, \mathsf{o_s}$ *are on the table:*

$$\phi, e, w, h \models [\mathsf{sense}][\mathsf{move}][\mathsf{sense}][\mathsf{fuse}]\mathbf{K}(x = \mathsf{o_m} \lor x = \mathsf{o_c} \lor x = \mathsf{o_s} \supset On(x))$$

Let z stand for $\langle \mathsf{sense}, \mathsf{move}, \mathsf{sense} \rangle$. For each $r \in \{\mathsf{o_m}, \mathsf{o_c}, \mathsf{o_s}\}$ we have $e, h_z^w, z \models \mathbf{S}^1_{\mathsf{sense}} On(r)$ and thus by condition (2), $w'[On(r), z \cdot \mathsf{fuse}] = 1$ for each $w' \in e{\downarrow}^{\phi, h^w_{z \cdot \mathsf{fuse}}, z \cdot \mathsf{fuse}}$.

4. Related Work

Reiter's situation calculus in its original form (Reiter, 2001) does not account for sensing actions and the agent's knowledge or belief. An epistemic extension by Scherl and Levesque (2003) added a possible worlds semantics within classical first-order logic. Lakemeyer and Levesque (2011) 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. This behaviour can be simulated in $\mathcal{ESF}$, as we have shown in Theorem 2.3. In that sense, $\mathcal{ESF}$ subsumes $\mathcal{ES}$ and, hence, also the Scherl–Levesque approach. A prominent feature of the Scherl–Levesque framework and $\mathcal{ES}$ is the extension of regression (Reiter, 2001) to the case of knowledge. The idea behind regression is to rewrite a query involving actions to an equivalent query over the initial situation only. That way, the projection problem is reduced to ordinary first-order reasoning without actions. Unfortunately, these results do not carry over to $\mathcal{ESF}$ in the general case due to our use of sensing histories. Alternatively, the projection problem can be solved through progression (Lin & Reiter, 1997), which modifies the knowledge base to account for action effects. How to compute progression in $\mathcal{ESF}$ is an open problem (cf. Section 5).

Other action formalisms like SADL (Golden & Weld, 1996), the event calculus (Forth & Shanahan, 2004), and the fluent calculus (Thielscher, 2000) have been 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, Halpern, and Levesque (1999) 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 (Bacchus et al., 1999), a variant by Belle and Levesque (2013) 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 (De Giacomo & Levesque, 1999) 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, Pagnucco, Lespérance, and Levesque (2011) presented a theory for belief change in the situation calculus. They use counterfactuals to specify an agent's preferential belief structure, and sensing results then trigger belief change. Counterfactuals like this could be applied to the cup versus mug problem: we believe that, if the object had a handle, it would be a mug. While Shapiro et al. (2011) assume sensing results to be correct, their work has recently been extended to also handle contradictory sensing results (Schwering, Lakemeyer, & Pagnucco, 2015), where strongest belief is always given to the most recent sensing. A combination of this framework and our approach of fusion actions might be interesting in order to achieve passive sensor fusion.

The closed-world assumption was introduced by Reiter (1978). Etzioni, Golden, and Weld (1994) 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. Whether the closed-world assumption is actually true – in which case Etzioni et al. refer to it as closed-world *information* – or not, is not explicit in our language. Several variants of the closed-world assumption have been proposed to overcome its limitations with disjunctive information, for example, (Minker, 1982).

The forgetting mechanism of $\mathcal{ESF}$ is essentially the same as the one proposed by Rajaratnam et al. (2014): it undoes the knowledge-effect of a previous action. This notion of forgetting is fundamentally different from Lin and Reiter's logical forgetting (Lin & Reiter, 1994). In particular, it is not possible to forget arbitrary facts or relations – unless just this fact or relation was learned through an epistemic action. We refer to (Rajaratnam et al., 2014) for the details.

KnowRob (Tenorth & Beetz, 2013) 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

The logic we presented in this paper aims at high-level sensor fusion, particularly to guide an active perception system, where an agent needs to *think* about how she could acquire new knowledge. This is a common problem in robotics, and our approach represents a possible solution from the KR perspective. In particular, our logic differs in the following ways from previous approaches:

- Actions may *sense an unbounded number of objects*, and the purported sensing results of an action may vary over the course of action. In particular, subsequent *sensings may complement or contradict each other.* We remark that this does not necessitate second-order logic. Which (perhaps incorrect) sensing results are returned is stipulated by the real world in our model.
- Sensing actions do not affect knowledge immediately. Instead, this effect is deferred to a *fusion action*, which produces new knowledge.
- Actions can apply a local *closed-world assumption* to solidify the agent's episteme.
- A *forgetting* mechanism allows to mitigate the epistemic effects of preceding actions.

The prominent role of active fusion, closed-world assumption, and forgetting may be considered as merely postponing a problem instead of solving it: whereas formalisms like the Scherl–Levesque framework require accurate modeling of sensing, $\mathcal{ESF}$ pushes that responsibility to the decision *when* to fuse, close, or forget. An interesting open problem is therefore to come up with a reasonable passive scheme that performs these epistemic operations when appropriate. One way to address this issue might be through ideas similar to iterated belief revision approaches, as they cope with inconsistent beliefs and sensings.

Another open problem is how to integrate numeric uncertainties reported by the sensors into the logic. Since we lack a probabilistic error model for the domains we have in mind (cf. Section 3.2), we are looking into possibilistic logic (Dubois & Prade, 2004) to address this issue.

Furthermore we are interested in an automated solution to the projection problem, such as regression. While regression appears to be infeasible in $\mathcal{ESF}$, the approach by Lakemeyer and Levesque (2014) may provide a solution, which is similar to progression. Based on this, we plan to deploy a decidable fragment of $\mathcal{ESF}$ in the spirit of (Lakemeyer & Levesque, 2014) on our robots. When dealing with infinite domains like the sonar distance in Section 3.1, it may be efficient to represent incomplete knowledge through intervals (Funge, 1999). The resulting system is intended as a reasoning back-end for high-level robot control programs written in Golog (H. Levesque, Reiter, Lésperance, Lin, & Scherl, 1997) in the context of our project on hybrid reasoning.[6]

References

Bacchus, F., Halpern, J. Y., & Levesque, H. J. (1999). Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, *111*(1–2), 171–208.

Bajcsy, R. (1988). Active perception. *Proceedings of the IEEE*, *76*(8), 966–1005.

Belle, V., & Levesque, H. J. (2013). Reasoning about continuous uncertainty in the situation calculus. In *Proceedings of the twenty-third international joint conference on artificial intelligence.*

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

De Giacomo, G., & Levesque, H. (1999). An incremental interpreter for high-level programs with sensing. In H. J. Levesque & F. Pirri (Eds.), *Logical foundations for cognitive agents: contributions in honor of Ray Reiter* (pp. 86–102). Springer.

Dubois, D., & Prade, H. (2004). Possibilistic logic: a retrospective and prospective view. *Fuzzy sets and systems*, *144*(1), 3–23.

Etzioni, O., Golden, K., & Weld, D. (1994). Tractable closed world reasoning with updates. In *Proceedings of the fourth international conference on principles of knowledge representation and reasoning (KR)* (pp. 178–189).

Forth, J., & Shanahan, M. (2004). Indirect and conditional sensing in the event calculus. In *Proceedings of the sixteenth European conference on artificial intelligence (ECAI)* (pp. 900–905).

Funge, J. (1999). Representing knowledge within the situation calculus using interval-valued epistemic fluents. *Reliable computing*, *5*(1), 35–61.

Golden, K., & Weld, D. (1996). Representing sensing actions: The middle ground revisited. In *Proceedings of the fifth international conference on principles of knowledge representation and reasoning (KR).*

Kripke, S. A. (1976). Is there a problem about substitutional quantification? In G. Evans & J. McDowell (Eds.), *Truth and meaning* (pp. 324–419). Oxford University Press.

Lakemeyer, G., & Levesque, H. J. (2004). Situations, si! situation terms, no! In *Proceedings of the ninth international conference on principles of knowledge representation and reasoning (KR)* (pp. 516–526).

Lakemeyer, G., & Levesque, H. J. (2011). A semantic characterization of a useful fragment of the situation calculus with knowledge. *Artificial Intelligence*, *175*(1), 142–164.

Lakemeyer, G., & Levesque, H. J. (2014). Decidable reasoning in a fragment of the epistemic situation calculus. In *Proceedings of the fourteenth international conference on principles of knowledge representation and reasoning (KR).*

Levesque, H., Reiter, R., Lésperance, Y., Lin, F., & Scherl, R. (1997). GOLOG: A logic programming language for dynamic domains. *Journal of Logic Programming*, *31*(1–3), 59–83.

Levesque, H. J., & Lakemeyer, G. (2001). *The logic of knowledge bases.* MIT Press.

Lin, F., & Reiter, R. (1994). Forget it! In *Proceedings of the AAAI fall symposium on relevance* (pp. 154–159).

Lin, F., & Reiter, R. (1997). How to progress a database. *Artificial Intelligence*, *92*(1), 131–167.

Minker, J. (1982). On indefinite databases and the closed world assumption. In D. Loveland (Ed.), *Sixth conference on automated deduction* (Vol. 138, pp. 292–308). Springer Berlin Heidelberg.

Niemueller, T., Abdo, N., Hertle, A., Lakemeyer, G., Burgard, W., & Nebel, B. (2013). Towards Deliberative Active Perception using Persistent Memory. In *Proceedings of the IROS 2013 workshop on AI-based robotics.*

Rajaratnam, D., Levesque, H. J., Pagnucco, M., & Thielscher, M. (2014). Forgetting in action. In *Proceedings of the fourteenth international conference on principles of knowledge representation and reasoning (KR).*

Reiter, R. (1978). On closed world data bases. In H. Gallaire & J. Minker (Eds.), *Logic and data bases* (pp. 55–76). Plenum Press.

Reiter, R. (2001). *Knowledge in action: Logical foundations for specifying and implementing dynamical systems.* The MIT Press.

Scherl, R., & Levesque, H. J. (2003). Knowledge, action, and the frame problem. *Artificial Intelligence*, *144*(1–2), 1–39.

Schwering, C., Lakemeyer, G., & Pagnucco, M. (2015). Belief revision and progression of knowledge bases in the epistemic situation calculus. In *Proceedings of the 24th international joint conference on artificial intelligence (IJCAI).*

Shapiro, S., Pagnucco, M., Lespérance, Y., & Levesque, H. J. (2011). Iterated belief change in the situation calculus. *Artificial Intelligence*, *175*(1), 165–192.

Tenorth, M., & Beetz, M. (2013). KnowRob – a knowledge processing infrastructure for cognition-enabled robots. Part 1: The knowrob system. *International Journal of Robotics Research (IJRR).*

Thielscher, M. (2000). Representing the knowledge of a robot. In *Proceedings of the seventh* *international conference on principles of knowledge representation and reasoning (KR)* (pp. 109–120).