Sampling-Based Belief Revision

Michael Thielscher
University of New South Wales, Australia
mit@unsw.edu.au

Abstract

Model sampling has proved to be a practically viable method for decision-making under uncertainty, for example in imperfect-information games with large state spaces. In this paper, we examine the logical foundations of sampling-based belief revision. We show that it satisfies six of the standard AGM postulates but not Vacuity nor Subexpansion. We provide a corresponding representation theorem that generalises the standard result from a single to a family of faithful assignments for a given belief set. We also provide a formal axiomatisation of sampling-based belief revision in the Situation Calculus as an alternative way of reasoning about actions, sensing, and beliefs.

1 Introduction

Model sampling is a viable approach to automated decision-making under uncertainty in a variety of domains; examples include the control of autonomous robots [Fox *et al.*, 1999; Hollinger and Sukhatme, 2014] and computer game playing, where sampling is the method of choice to cope with highly incomplete knowledge of the game state in most imperfect-information games [Frank and Basin, 1998; Ginsberg, 2001; Richards and Amir, 2012]. In all of these applications, agents base their decisions on a small set of samples from the set of all complete models that are consistent with their observations. When further information is acquired, models that turn out to be inconsistent with the new observations are *resampled* in order to maintain both consistency as well as a uniform size of the sample set [Frank and Basin, 1998; Edelkamp *et al.*, 2012; Schofield *et al.*, 2012].

From a logical perspective, when agents base their decisions on model sampling then they can be said to *believe* that the world is fully characterised by the samples. If, for instance, all models randomly sampled by a Bridge-playing agent [Ginsberg, 2001] happen to assume that their partner carries at least one card of each suit, then the agent concludes that the player will have to "follow suit" on any lead card. However, rational agents operating with beliefs (rather than knowledge) ought to take into account the possibility that their beliefs may be wrong, for example when a subsequent observation implies that a player does in fact not have a card of a specific suit. When this happens, the agent needs to *resample* inconsistent models [Frank and Basin, 1998], thereby causing him to *revise* his beliefs in the light of new information.

In this paper, we study the formal properties of this kind of belief revision operators that follow from the sampling-based approach to handling incomplete information. The purpose of this study is two-fold:

1. A formal analysis of how model sampling affects the beliefs of an agent provides for a deeper understanding of the formal properties of an approach that has proved successful in a range of applications including robot control and computer game playing.
2. A formalisation of model sampling as a general principle for belief revision enriches the suite of traditional operators based on the AGM postulates [Alchourrón *et al.*, 1985], especially since it promises to be a viable method for applications with large state spaces for which classical Belief Revision operators are not practical.

Our investigation into the logical foundations for sampling-based belief revision will show that it satisfies six of the eight standard AGM postulates but not *Vacuity* nor *Subexpansion*. The intuitive reason is that when some but not all sampled models are consistent with new observation ϕ, then those that contradict ϕ get resampled. These new models may then invalidate previously held beliefs despite the fact that the presence of other samples that are consistent with ϕ imply that the new observation is consistent with the agent's original beliefs.

As a simple example, consider a situation in an unspecified card game where there are the 6 outstanding cards $\clubsuit 2, \spadesuit 2, \clubsuit 3, \heartsuit 3, \diamondsuit 3, \diamondsuit 4$, and suppose that an agent has randomly generated the following two samples for the belief about the remaining hand (consisting of two cards) of another player: $K = \{\{\clubsuit 3, \diamondsuit 3\}, \{\clubsuit 2, \heartsuit 3\}\}$. Suppose further that the player is observed not to follow suit when hearts is played. This observation, which we shall write $\phi = \neg\heartsuit$, renders the second model inconsistent (but not the first). If we assume that resampling this model results in the new sample set $K * \phi = \{\{\clubsuit 3, \diamondsuit 3\}, \{\spadesuit 2, \diamondsuit 4\}\}$, then the following can be said about the beliefs of the agent before and after observing ϕ:

1. In K, the agent believes that the player has a $\clubsuit$-card.
2. K is consistent with $\phi = \neg\heartsuit$.

(AGM-1)	$K * \phi = Cn(K * \phi)$	(Closure)
(AGM-2)	$\phi \in K * \phi$	(Success)
(AGM-3)	$K * \phi \subseteq K + \phi$	(Inclusion)
(AGM-4)	If $\neg\phi \notin K$ then $K * \phi \supseteq K + \phi$	(Vacuity)
(AGM-5)	If ϕ is consistent, then $K * \phi$ is consistent	(Consistency)
(AGM-6)	If $\phi \equiv \psi$, then $K * \phi = K * \psi$	(Extensionality)
(AGM-7)	$K * (\phi \wedge \psi) \subseteq (K * \phi) + \psi$	(Superexpansion)
(AGM-8)	If $\neg\psi \notin (K * \phi)$, then $K * (\phi \wedge \psi) \supseteq (K * \phi) + \psi$	(Subexpansion)

Figure 1: The eight AGM postulates.

3. In $K * \phi$, the agent no longer believes that the player has a ♣-card.

These three properties together violate the Vacuity postulate, which satipulates that if ϕ is consistent with K, then every belief in K should also be a belief in $K * \phi$. We will show that for a similar reason, the Subexpansion postulate is not satisfied either.

We will provide a *representation theorem* for belief revision operators that satisfy all but these two AGM postulates. The result will be a generalisation of the classical representation theorem [Gärdenfors and Makinson, 1988] from single to a family of faithful assignments for a given belief set.

The remainder of the paper is organised as follows. In the next section, we recall the necessary notions and notations for belief revision and logical reasoning about actions. In the section that follows, we define and analyse model sampling for belief revision in the context of the AGM framework. Thereafter, we state and prove the representation result. Finally, we apply our general concept of model sampling for belief revision to develop a variant of the classical Situation Calculus for reasoning about actions, sensing and beliefs based on (re-)sampling.

2 Background

We will formulate and analyse sampling-based belief revision in the context of the classical AGM-based approach and assuming a logic language $\mathcal{L}$ generated from a finite set $\mathcal{P}$ of propositions. By $Cn(\phi)$ we denote all classical logical consequences of a sentence, or a set of sentences, ϕ. $K + \phi$ stands for $Cn(K \cup \{\phi\})$. As usual, a propositional interpretation (world) is a mapping from $\mathcal{P}$ to $\{\top, \bot\}$. The set of all interpretations is denoted by $\mathcal{W}$. If an interpretation w truth-functionally maps a sentence, or a set of sentences, ϕ to $\top$, then w is called a model of ϕ (denoted by $w \models \phi$). $Mods(\phi)$ denotes the set of all models of ϕ, and the latter is consistent iff $Mods(\phi) \neq \emptyset$. Formulas, or sets of formulas, ϕ and ψ are equivalent, written $\phi \equiv \psi$, iff $Mods(\phi) = Mods(\psi)$.

AGM Postulates Let the beliefs of an agent be represented by a set of sentences K in $\mathcal{L}$. Any new evidence is a sentence ϕ in $\mathcal{L}$, and the result of *revising* K with ϕ, denoted by $K * \phi$, is also a belief set over $\mathcal{L}$. Alchourrón *et al.* [1985] put forward eight postulates for belief revision operators, which are given in Figure 1. Readers are referred to Gärdenfors and Makinson [1988] for a detailed discussion of the motivation and interpretation of these postulates.

A Representation Theorem The classical representation theorem [Gärdenfors and Makinson, 1988] characterises belief revision operators that satisfy the AGM postulates with the help of preference relations over worlds (i.e. models). A total pre-order $\leq$ (possibly indexed) is a reflexive, transitive binary relation such that either $\alpha \leq \beta$ or $\beta \leq \alpha$ holds for any α, β. The strict part of $\leq$ is denoted by $<$, that is, $\alpha < \beta$ iff $\alpha \leq \beta$ and $\beta \not\leq \alpha$. As usual, $\alpha = \beta$ abbreviates $\alpha \leq \beta$ and $\beta \leq \alpha$. Given any total pre-order $\leq$ over a set S, by $\min(S, \leq)$ we denote the set $\{\alpha \in S : \alpha \leq \beta \text{ for all } \beta \in S\}$.

Definition 1 *A function that maps each belief set K to a total pre-order $\leq_K$ on $\mathcal{W}$ is called a* faithful assignment *if it satisfies the following conditions.*

- *If $w_1, w_2 \models K$, then $w_1 =_K w_2$.*
- *If $w_1 \models K$ and $w_2 \not\models K$, then $w_1 <_K w_2$.*
- *If $J \equiv K$, then $\leq_J = \leq_K$.*

The intuitive meaning of $w_1 \leq_K w_2$ is that w_1 is at least as plausible as w_2 from the viewpoint of the agent who possesses the belief set K. The classical representation theorem says that belief revision operators that satisfy the AGM postulates can be characterised with the help of faithful assignments and vice versa.

Theorem 1 *A revision operator $*$ satisfies Postulates (AGM-1)–(AGM-8) iff there is a faithful assignment that maps a belief set K to a total pre-order $\leq_K$ such that*

$$\text{Mods}(K * \phi) = \min(\text{Mods}(\phi), \leq_K)$$

Reasoning About Actions Model-sampling has been successfully used as a general technique for agents that act and sense in a dynamic environment which they cannot fully observe. This includes the control of autonomous robots [Fox *et al.*, 1999; Hollinger and Sukhatme, 2014] and computer players for imperfect-information games [Frank and Basin, 1998; Ginsberg, 2001; Richards and Amir, 2012]. Action calculi for knowledge representation provide the logical foundations for reasoning about actions, with the Situation Calculus being the classical one [McCarthy, 1963]. We briefly recapitulate the basic concepts of this calculus in a variant that uses a special fluent to represent the knowledge of agents [Scherl and Levesque, 2003].The Situation Calculus is a predicate logic with a few pre-defined language elements:

- constant s_0, which denotes the *initial situation*, along with constructor $\textsc{Do}(A, S)$ to denote the situation resulting from doing *action* A in situation S;

- predicate $\text{HOLDS}(F, S)$, which denotes that *fluent* F (i.e., an atomic state feature) is true in situation S;
- predicate $\text{POSS}(A, S)$, which denotes that action A is possible in situation S.

For example, a fluent $\mathit{HasCard}(P, C, V)$ may be used to indicate that player P holds a card of suit C and value V. The effect of playing a specific card, represented by the action $\mathit{Play}(P, C, V)$, on all cards being held could then be axiomatised by the following so-called *successor state axiom* [Reiter, 1991] (free variables are assumed to be universally quantified):

$$\begin{aligned}&\text{POSS}(A, S) \rightarrow \\ &\quad (\,\text{HOLDS}(\mathit{HasCard}(P, C, V), \text{DO}(A, S)) \leftrightarrow \\ &\qquad \text{HOLDS}(\mathit{HasCard}(P, C, V), S) \wedge A \neq \mathit{Play}(P, C, V)\,)\end{aligned}$$

Put in words, all cards held in situation S will also be held in the situation after doing action A in S except for the card that happens to have been played.

Scherl and Levesque [2003] use the special fluent $\mathbf{K}(S', S)$—to be read as: situation S' is accessible from situation S—in order to formally represent, and reason about, the knowledge of an agent within the Situation Calculus:

$$\text{KNOWS}(\Phi, S) \stackrel{\text{def}}{=} \forall S'. \mathbf{K}(S', S) \rightarrow \Phi[S'] \tag{1}$$

Here, Φ is a *reified* formula in which fluents are used as predicates (without situation argument); and $\Phi[S']$ is Φ with all fluents F replaced by $\text{HOLDS}(F, S')$. For example,

$$\text{KNOWS}(\forall X. \neg \mathit{HasCard}(p, \heartsuit, X), s_0)$$

is $\forall S'. \mathbf{K}(S', s_0) \rightarrow \forall X. \neg \text{HOLDS}(\mathit{HasCard}(p, \heartsuit, X), S')$, meaning that the agent knows player p has no $\heartsuit$-card in the initial situation as this is true for all s_0-accessible situations.

The effects of actions, including sensing, on the knowledge state of an agent are defined by the successor state axiom for the special fluent **K** [Scherl and Levesque, 2003]:

$$\begin{aligned}&\text{POSS}(A, S) \rightarrow \\ &\quad \mathbf{K}(S'', \text{DO}(A, S)) \leftrightarrow \\ &\qquad \exists A', S'. S'' = \text{DO}(A', S') \wedge \mathbf{K}(S', S) \wedge \\ &\qquad\qquad \forall P. \text{SENSE}(P, A, S) \leftrightarrow \text{SENSE}(P, A', S')\end{aligned} \tag{2}$$

Put in words, S'' is a possible situation after action A in S if, and only if, S'' is obtained by doing A in a situation S' that was conceivable in S; and the agent's sensing result for A', S' is equivalent to his sensing result for the actual A, S.[1]

3 Model Sampling for Belief Revision

In sampling-based approaches to decision making under uncertainty, agents approximate their incomplete knowledge about the state of a system with the help of a fixed, and typically small, number of sampled models that are consistent with given observations. At any point in time, the current collection of samples can be viewed as the agent's belief about the current state. New observations act as filters through which only those samples are kept that are consistent with the new information. Since the accuracy of the belief usually decreases with a shrinking number of samples, it is important that models which have been filtered out are being resampled, that is, get replaced by new and consistent samples. The belief of the agent thus gets revised. In the following, we will formalise this very general concept of (re-)sampling in the context of belief revision in order to analyse its underlying properties.

Various strategies have been used for the actual process of generating a set of models that are consistent with given observations. They vary from be highly domain-specific [Ginsberg, 2001] to domain-independent but according to a specific process [Richards and Amir, 2012; Schofield *et al.*, 2012] to purely random resampling [Fox *et al.*, 1999]. To accommodate these different types of strategies in our systematic analysis of how resampling changes the beliefs of an agent, we assume that models are generated in an arbitrary but fixed order, which may have been obtained by a purely random process.

Example 1 *Suppose the beliefs about the remaining two cards of a player are generated from 6 outstanding cards according to some, possibly purely random, sampling strategy in the order shown in Figure 2. Let us assume an agent who maintains a sample set S of size 2, then the initial set of samples—without further information—would consist of the first two in the series of generated models, that is,*

$$S = \{w_1, w_2\} = \{\{\clubsuit 3, \diamondsuit 3\}, \{\clubsuit 2, \heartsuit 3\}\} \tag{3}$$

Now, consider two examples of new observations.

First, suppose we learn that the player whose hand we sample does not have a $\heartsuit$-card. The first model, w_1, is consistent with this new information and hence will be retained, whereas the second sample, w_2, is not. Resampling the inconsistent model means to consider the next world in the sequence of generated samples that is consistent with $\neg\heartsuit$, which is w_4. This results in the new sample set[2]

$$S \triangleleft (\neg\heartsuit) = \{w_1, w_4\} = \{\{\clubsuit 3, \diamondsuit 3\}, \{\spadesuit 2, \diamondsuit 4\}\} \tag{4}$$

Second, suppose we make the more specific observation that the player does not own a $\heartsuit$-card and *that one of his cards is $\diamondsuit 3$. Again, w_1 is consistent and remains in the sample set while the second sample is replaced by the next model in sequence after w_2 and consistent with the new information $\neg\heartsuit \wedge \diamondsuit 3$, which is w_6. This results in the new sample set*

$$S \triangleleft (\neg\heartsuit \wedge \diamondsuit 3) = \{w_1, w_6\} = \{\{\clubsuit 3, \diamondsuit 3\}, \{\diamondsuit 3, \diamondsuit 4\}\} \tag{5}$$

The *beliefs* of an agent who approximates the state of a system by a set of samples are all the logical consequences of the sample models.

Definition 2 *The* belief set *determined by a finite set of samples $S = \{w_1, \dots, w_n\} \subseteq \mathcal{W}$ is given by*

$$K|_S = \{\alpha : w_i \models \alpha \text{ for all } w_i \in S\} \tag{6}$$

Recall, for instance, the initial sample set S given by (3) in Example 1. It follows that $\neg\heartsuit \notin K|_S$ but $(\clubsuit 2 \vee \clubsuit 3) \in K|_S$.

We are now in a position to formalise the concept of resampling in belief revision, based on an arbitrary but fixed strategy for the generation of sample models.

[1] $\text{SENSE}(P, A, S)$ means that P, the encoding of a potential sensing result, is indeed perceived after doing action A in situation S.

[2] The notation $S \triangleleft \phi$ below is used to denote the result of resampling a set of models S upon making a new observation expressed by sentence ϕ.

$\{\clubsuit 3, \diamondsuit 3\}$	$\{\clubsuit 2, \heartsuit 3\}$	$\{\heartsuit 3, \diamondsuit 4\}$	$\{\spadesuit 2, \diamondsuit 4\}$	$\{\spadesuit 2, \heartsuit 3\}$	$\{\diamondsuit 3, \diamondsuit 4\}$	$\dots$	$\{\heartsuit 3, \diamondsuit 3\}$
w_1	w_2	w_3	w_4	w_5	w_6	$\dots$	w_{15}

Figure 2: A random sequence of samples for a 2-card hand generated from the 6 outstanding cards $\clubsuit 2, \spadesuit 2, \clubsuit 3, \heartsuit 3, \diamondsuit 3, \diamondsuit 4$.

Definition 3 *Let S be a given sample set of size $n > 0$. Consider an arbitrary but fixed order $w_1 <_S w_2 <_S w_3 <_S \dots$ on $\mathcal{W} \setminus S$ in which further samples are generated. Let ϕ be a sentence, then* resampling S for ϕ *results in the set of worlds*

$$S \lhd \phi \;=\; (S \cap \text{Mods}(\phi)) \,\cup\, \min_k(\text{Mods}(\phi), <_S) \tag{7}$$

where

- $k = n - |S \cap \text{Mods}(\phi)|$
- $\min_k(\text{Mods}(\phi), <_S)$ *are the m smallest worlds in $w_1 <_S w_2 <_S w_3 < \dots$ that are models of ϕ but not in S, where $m = \min(k, |\text{Mods}(\phi)|)$.*[3]

Put in words, all samples consistent with ϕ are retained, and all other models are replaced by the next m generated worlds that are consistent with ϕ and different from any of the retained samples. If ϕ admits a sufficient number of models, then the sample size n is retained (case $m = k$), otherwise it is reduced to the number of remaining consistent models.

Example 2 *Recall Example 1. It is easy to see that for sample size $n = 2$, (4) and (5) satisfy the conditions of Definition 3 wrt. the sampling order given in Figure 2. It is also easy to see that there is only one world that satisfies $\neg\clubsuit \wedge \neg\diamondsuit$, hence $S \lhd (\neg\clubsuit \wedge \neg\diamondsuit) = \{\{\spadesuit 2, \heartsuit 3\}\}$ must result in a set smaller than the original sample size of 2. Resampling for an observation that has no model results in the empty set, e.g. $S \lhd (\clubsuit 3 \wedge \neg\clubsuit 3) = \emptyset$.*

As the main result in this section, we show that sampling-based belief revision satisfies six of the eight classical AGM postulates but not Vacuity nor Subexpansion (cf. Figure 1).

Theorem 2 *Let $K|_S$ be a belief set determined by a sample set S of size $n > 0$ and $\lhd$ the operation of resampling wrt. any given order for generating samples. The corresponding belief revision operator $*$ defined by*

$$K|_S * \phi = K|_{S \lhd \phi} \tag{8}$$

*satisfies (*AGM-*1)–(*AGM-*3) and (*AGM-*5)–(*AGM-*7).*

Proof: We prove each postulate in turn.

- (AGM-1) follows by Definition 2, equation (6).
- (AGM-2) holds by Definition 3, equation (7), which implies that $S \lhd \phi \subseteq \mathit{Mods}(\phi)$, hence $\phi \in K|_{S \lhd \phi}$ according to equation (6).
- For (AGM-3) we observe that if $\alpha \in K|_S * \phi$ then $w \models \alpha$ for all $w \in S \lhd \phi$ according to (8), hence $w \models \alpha$ for all $w \in S \cap \mathit{Mods}(\phi)$ according to (7). It follows that $\alpha \in K|_S + \phi$ according to (6).
- (AGM-5) holds since $\mathit{Mods}(\phi) \neq \emptyset$ implies $S \lhd \phi \neq \emptyset$ according to (7); hence, if ϕ is consistent then by (8), $K|_S * \phi$ is consistent too.
- (AGM-6) follows from Definition 3, equation (7), since $\phi \equiv \psi$ implies $\mathit{Mods}(\phi) = \mathit{Mods}(\psi)$, hence $S \lhd \phi = S \lhd \psi$, which implies $K|_S * \phi = K|_S * \psi$ by (8).
- For (AGM-7), recall Definition 3, equation (7), and consider any sample $w \in S \lhd \phi$ that satisfies not only ϕ but also ψ. By Definition 3 and because the same sampling strategy $<_S$ is used when constructing $S \lhd (\phi \wedge \psi)$, it follows that w is either among the samples in S that satisfy $\phi \wedge \psi$, or is among the first m generated samples that satisfy $\phi \wedge \psi$, where m is given as in Definition 3. Hence, $w \in S \lhd (\phi \wedge \psi)$. Therefore, $(S \lhd \phi) \cap \mathit{Mods}(\psi) \subseteq S \lhd (\phi \wedge \psi)$. This implies that $K|_{S \lhd \phi} + \psi \supseteq K|_{S \lhd (\phi \wedge \psi)}$, hence $K|_S * (\phi \wedge \psi) \subseteq (K|_S * \phi) + \psi$ according to (8).

$\square$

Notably, the fact that sampling-based belief revision satisfies (AGM-3) and (AGM-7) guarantees that arbitrary new beliefs, i.e. which do not follow from $K + \phi$ or $(K * \phi) + \psi$, are not introduced when revising a sample set. On the other hand, it is easy to construct counterexamples to show that neither Vacuity (AGM-4) nor Subexpansion (AGM-8) are satisfied.

Example 3 *Recall S as defined in (3). For Vacuity, note, e.g., that $(\clubsuit 2 \vee \clubsuit 3) \in K|_S$ and that $\phi = \neg\heartsuit$ is consistent with $K|_S$. However, according to (4), $(\clubsuit 2 \vee \clubsuit 3) \notin K|_S * \phi$.*

For Subexpansion, observe first that $\diamondsuit 3$ is consistent with $S \lhd \neg\heartsuit$ according to (4), hence $\neg\diamondsuit 3 \notin K * (\neg\heartsuit)$. From (4) we can also see that $\clubsuit 3 \in (K * (\neg\heartsuit)) + \diamondsuit 3$. However, (5) implies that $\clubsuit 3 \notin K * (\neg\heartsuit \wedge \diamondsuit 3)$.

4 A Representation Theorem

The classical representation theorem for AGM-based belief revision states that any revision operator that satisfies all classical postulates can be equivalently characterised as faithful assignments of every belief set to a specific preference relation over worlds (cf. Definition 1 and Theorem 1). We have seen that sampling-based belief revision satisfies all but two of the postulates. This begs the question whether the classical representation theorem can be generalised so as to characterise belief revision operators that comply with the standard postulates minus Vacuity and Subexpansion.

It turns out that indeed a suitably extended representation theorem can be given based on the same concept of faithful assignments and preference orderings. The necessary weakening of the conditions of Theorem 1 is obtained by allowing for *multiple* total pre-orders for a given belief set. This can be intuitively explained as follows: Different possible worlds may be characterised by different underlying assumptions, which then lead to different preference orderings among alternative worlds upon revision. When revising a set of preferred worlds in the light of a family of faithful assignments, the resulting set is obtained as the combined set of minimal worlds from all preference relations.

[3] If $|Mods(\phi)| = 0$ then $S = \emptyset$, i.e., no consistent sample exists.

Definition 4 *A mapping from belief sets K to collections of total pre-orders $\leq_K^1, \ldots, \leq_K^n$ ($n \geq 1$) on $\mathcal{W}$ is a* weakly faithful assignment *if it satisfies the following conditions.*

- *If $w_1 \models K$ then there is a $k \in \{1, \ldots, n\}$ such that $w_1 \leq_K^k w_2$ for all w_2.*
- *If $w_2 \not\models K$ then for all $k \in \{1, \ldots, n\}$ there is some $w_1 \models K$ such that $w_1 <_K^k w_2$.*
- *If $J \equiv K$, then $\leq_J^1, \ldots, \leq_J^n = \leq_K^1, \ldots, \leq_K^n$.*

Put in words, every K-world is most preferred in at least one ordering (but may be even less preferred than a non-K world in others) while non-K worlds are never *most* preferred. It is easy to verify that all faithful assignments according to Definition 1 are also weakly faithful (with $n = 1$). As a generalisation of the classical representation theorem, revision with a weakly faithful assignment can be shown to characterise exactly the belief revision operators that satisfy all AGM-postulates except Vacuity and Subexpansion.

Theorem 3 *A revision operator $*$ satisfies Postulates (AGM-1)–(AGM-3) and (AGM-5)–(AGM-7) iff there is a weakly faithful assignment $\leq_K^1, \ldots, \leq_K^n$ such that*

$$\text{Mods}(K * \phi) = \bigcup_k \min(\text{Mods}(\phi), \leq_K^k) \tag{9}$$

Proof (sketch): The "only-if" direction can be proved by constructing, for a given $*$ and K, assignments $\leq_K^i$, one for each $w_i \in \mathit{Mods}(K)$, as follows:

- $w_i <_K^i w'$ for all $w' \in \mathcal{W} \setminus \{w_i\}$.
- $w' \leq_K^i w''$ if $w' \in \mathit{Mods}(K * \mathit{form}(w', w''))$,[4] for all $w', w'' \in \mathcal{W} \setminus \{w_i\}$.

(We note that this shows how to determine a corresponding system of spheres for an operator obtained from a given ordering of samples as per Definition 3.) Proving this to be weakly faithful is akin to the proof for a similarly constructed single assignment to be faithful in case of classical belief revision [Katsuno and Mendelzon, 1992].

For the "if" direction, we assume a revision operator given by (9) and prove that it satisfies each postulate in turn.

(AGM-1) follows from the definition of $K * \phi$ as the set of formulas whose models are given by (9).

(AGM-2) follows from $\mathit{Mods}(K * \phi) \subseteq \mathit{Mods}(\phi)$ by (9).

(AGM-3): If $w \in \mathit{Mods}(K \cup \{\phi\})$ then $w \in \mathit{Mods}(K) \cap \mathit{Mods}(\phi)$, hence $w \in \min(\mathit{Mods}(K), \leq_K^k) \cap \mathit{Mods}(\phi)$ for some $k \in \{1, \ldots, n\}$ by the first condition in Definition 4. Therefore, $\bigcup_k \min(\mathit{Mods}(\phi), \leq_K^k) \supseteq \mathit{Mods}(K \cup \{\phi\})$, which implies $K * \phi \subseteq K + \phi$.

(AGM-5) holds since consistency of ϕ, i.e. $\mathit{Mods}(\phi) \neq \emptyset$, implies that $\min(\mathit{Mods}(\phi), \leq_K^k) \neq \emptyset$ for all $1 \leq k \leq n$, hence $\mathit{Mods}(K * \phi) \neq \emptyset$.

(AGM-6) follows from the fact that whenever $\phi \equiv \psi$ then $\mathit{Mods}(\phi) = \mathit{Mods}(\psi)$, which for all $1 \leq k \leq n$ implies $\min(\mathit{Mods}(\phi), \leq_K^k) = \min(\mathit{Mods}(\psi), \leq_K^k)$.

(AGM-7): To show that $\bigcup_k \min(\mathit{Mods}(\phi), \leq_K^k) \cap \mathit{Mods}(\psi)$ is a subset of $\bigcup_k \min(\mathit{Mods}(\phi \wedge \psi), \leq_K^k)$, consider an arbitrary world $w \in \bigcup_k \min(\mathit{Mods}(\phi), \leq_K^k) \cap \mathit{Mods}(\psi)$, then $w \in \mathit{Mods}(\phi \wedge \psi)$. Suppose that for this world we have that $w \notin \bigcup_k \min(\mathit{Mods}(\phi \wedge \psi), \leq_K^k)$, then for all $1 \leq k \leq n$ there must be a world $w'_k \in \mathit{Mods}(\phi \wedge \psi)$ such that $w'_k <_K^k w$, contradicting the assumption $w \in \bigcup_k \min(\mathit{Mods}(\phi), \leq_K^k)$. □

[4] $\mathit{form}(w', w'')$ denotes an arbitrary propositional formula whose only models are w' and w'' [Katsuno and Mendelzon, 1992].

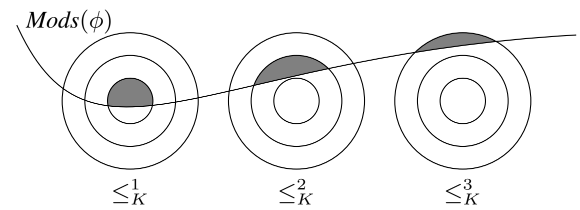

Figure 3: Revision with multiple systems of spheres [Grove, 1988]: If the three inner circles together are the models for the current belief, K, then the worlds in the shaded areas together are the models for the revised belief set $K * \phi$.

The principle of combining the minimal ϕ-models from multiple systems of spheres is graphically illustrated in Figure 3: Even in case some K-worlds are consistent with the new observation (here, some of the most preferred worlds in $\leq_K^1$), the revised set may include new models determined by an alternative preference relation (here, $\leq_K^2$ and $\leq_K^3$). This is the defining characteristic of operators that do not satisfy Vacuity or Subexpansion but all other AGM postulates.

A justification for the use of more than one faithful assignment can be that a belief may have been formed by considering alternative hypotheses, as in the following example.

Example 4 *We know that when John goes out for dinner, he always chooses between Asian (A) and Western (W) cuisine. He prefers to have the former in Thai Town (T) and the latter elsewhere ($\neg T$). We also know that he usually goes for dessert (D) in a Western but not in an Asian restaurant. There is a notable exception with one of his favourite Asian restaurants, however, where he usually does have dessert. This restaurant happens not to be in Thai Town.*

Suppose John tells us that he went out for dinner last night, then we may represent uncertainty about his choice of cuisine by considering two hypotheses based on his preferred choices ($A \vee W$): on the one hand, if *he decided in favour of Asian food, then he would have gone to Thai Town and skipped dessert, so we believe in the conditional $A \to T \wedge \neg D$; on the other hand, he could have had Western food elsewhere, including dessert, so we also believe in the conditional $W \to \neg T \wedge D$. Altogether our belief is characterised by*

$$(A \wedge T \wedge \neg D) \vee (W \wedge \neg T \wedge D) \tag{10}$$

Now, suppose we subsequently learn that John had in fact dessert, then we may reconsider both principled hypotheses in turn: while the second (W) is consistent with the additional observation, the first (A) is not. Rather than dropping this hypothesis, however, we may choose to select the closest world according to John's preferences in case of Asian cuisine, which means to consider the possibility that he went to his preferred Asian restaurant that is not in Thai Town where he usually has dessert. Our revised belief will be $(A \wedge \neg T \wedge D) \vee (W \wedge \neg T \wedge D)$, which entails $\neg T \wedge D$ but not W, contrary to Vacuity applied to $(10) \in K$, $\phi = D$.

5 Axiomatising Model Sampling for Reasoning About Actions and Sensing

An interesting application of sampling-based logical reasoning about beliefs and how to revise them can be found in general theories of actions and sensing, aiming at a practically viable method to handle large domains in which agents have highly incomplete knowledge. A point in case are existing implementations of the Situation Calculus with sensing that attempt to maintain a complete description of the knowledge state, such as Reiter's [2001] use of prime implicants, which do not scale well due to the high theoretical complexity of executing knowledge-based plans [Lang and Zanuttini, 2012].

In the following, we develop a formal axiomatisation of model sampling in the Situation Calculus as an alternative way of reasoning about knowledge and akin to the modelling of classical belief revision in this calculus [Shapiro *et al.*, 2000]. To this end, we replace Scherl and Levesque's [2003] special epistemic fluent $\mathbf{K}(S', S)$ (to express that situation S' is accessible from situation S) by an alternative fluent $\mathbf{S}(S', S)$ whose intuitive meaning is that of situation S' being a *sample* obtained for the actual situation S.

It is of course not always the case that the actual situation is among the sampled ones. In fact this is very unlikely in practice. Hence, the axiom of reflexivity $\forall S.\, \mathbf{K}(S, S)$, which is a foundational axiom in the Situation Calculus with knowledge [Scherl and Levesque, 2003], does *not* translate into an axiom $\forall S.\, \mathbf{S}(S, S)$. As a consequence, unlike fluent $\mathbf{K}$ which axiomatises knowledge (cf. (1)), fluent $\mathbf{S}$ axiomatises belief. This is reflected in the following definition:

$$\textsc{Bel}(\Phi, S) \stackrel{\text{def}}{=} \forall S'.\, \mathbf{S}(S', S) \rightarrow \Phi[S'] \tag{11}$$

The adaptation of the successor state axiom for the knowledge fluent (cf. (2)) to the sampling fluent requires to axiomatise the concept of *resampling*. This can be formally captured by an auxiliary fluent $\mathbf{SR}(S'', S', T)$ meaning that situation S'' is a sample of situation S' *obtained through resampling in situation* T. The third argument is necessary because, for example, a situation s_0' can be a valid sample for s_0 when the agent is in that situation, but after sensing the agent may know enough to rule out that he could have been in s_0' initially. The following axioms for $\mathbf{S}$ and $\mathbf{SR}$ formalise the relation among samples between successor situations:

$$\begin{array}{l}
\textsc{Poss}(A, S) \rightarrow \\
\quad \mathbf{S}(S'', \textsc{Do}(A, S)) \rightarrow \\
\qquad \exists A', S'.\, S'' = \textsc{Do}(A', S') \wedge \mathbf{S}(S', S) \wedge \\
\qquad\qquad\quad \forall P.\, \textsc{Sense}(P, A, S) \leftrightarrow \textsc{Sense}(P, A', S') \\
\qquad \vee\, \mathbf{SR}(S'', \textsc{Do}(A', S'), \textsc{Do}(A', S')) \\
\\
\textsc{Poss}(A, S) \rightarrow \\
\quad \mathbf{SR}(S'', \textsc{Do}(A, S), T) \rightarrow \\
\qquad \exists A', S'.\, S'' = \textsc{Do}(A', S') \wedge \mathbf{SR}(S', S, T) \wedge \\
\qquad\qquad\quad \forall P.\, \textsc{Sense}(P, A, S) \leftrightarrow \textsc{Sense}(P, A', S')
\end{array} \tag{12}$$

In words, a sample S'' must be

- the result of updating a sample S' from the previous situation such that the sensing result coincides with what is sensed in the actual situation S;
- or the result of re-sampling in the resulting situation $\textsc{Do}(A, S)$.

Moreover, in order for a situation S'' to be the result of resampling situation $\textsc{Do}(A, S)$ in situation T, S'' must be obtained from a sample S' of S obtained through resampling in the same situation T.

We note that axioms (12) do not account for a specific sampling strategy; they rather define the range of all possible ways in which an agent can sample situations. Hence, (12) provides a mere *schema* from which actual successor state axioms (with the usual bi-implication [Reiter, 1991]) can be obtained by strengthening each right-hand side with a formula characterising a particular resampling strategy. This may include a limit on the sample size for each situation, e.g. to just two via the axiom $\mathbf{S}(S_1, S) \wedge \mathbf{S}(S_2, S) \wedge \mathbf{S}(S_3, S) \rightarrow S_1 = S_2 \vee S_1 = S_3 \vee S_2 = S_3$.

6 Conclusion

Motivated by the success of sampling technique in different domains in which agents have to reason and make decisions under uncertainty, we have examined the logical foundations of sampling-based belief revision. It turned out that operators based on this principle satisfy six of the eight standard AGM postulates but not Vacuity nor Subexpansion. The reason is that resampling may result in new models that invalidate previously held beliefs even in cases where a new observation is consistent with the agent's original beliefs. Hence, even if new information is consistent with currently held beliefs, nonetheless some of these beliefs may get dropped. But our analysis also shows that arbitrary new beliefs (i.e., which do not follow from the previously held beliefs and the new information) can never be introduced through model resampling. A corresponding representation theorem adapts the classical result from single to multiple systems of spheres for a given belief set. As an application of our general framework for sampling-based belief revision, we have adapted a formal Situation Calculus axiomatisation to provide an alternative way of reasoning about actions and beliefs using (re-)sampling.

A practically viable method for belief revision, model sampling enriches the suite of traditional operators based on the AGM postulates, especially for domains where maintaining a full set of models (explicitly or implicitly represented) is practically impossible due to large information sets. These properties provide important insights into the side-effects of resampling on the beliefs of a decision-making agent that reasons based on sampled models. This insight follows from the analysis of resampling as single-step revision, which covers the essence of the operation performed by the agent when processing new information. Nonetheless for future work it would be interesting to formally investigate iterated revision in this context [Darwiche and Pearl, 1997]. We are also interested in looking at other applications of belief revision in which the sampling-based technique may provide an interesting alternative to classical operators. Another question for future work is to analyse the formal relation between our axiomatisation of sampling-based reasoning in the Situation Calculus and the existing axiomatisation of AGM-style belief revision in this calculus [Shapiro *et al.*, 2000].

Acknowledgments

We thank the reviewers of an earlier version for their very helpful comments and suggestions. This research was supported by the Australian Research Council under grant no. DP150101351. The author is also affiliated with the University of Western Sydney.

References

[Alchourrón *et al.*, 1985] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[Darwiche and Pearl, 1997] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.

[Edelkamp *et al.*, 2012] Stefan Edelkamp, Tim Federholzner, and Peter Kissmann. Searching with partial belief states in general games with incomplete information. In B. Glimm and A. Krüger, editors, *Proceedings of the German Annual Conference on Artificial Intelligence (KI)*, volume 7526 of *LNCS*, pages 25–36, Saarbrücken, Germany, September 2012. Springer.

[Fox *et al.*, 1999] Dieter Fox, Wolfram Burgard, Frank Dellaert, and Sebastian Thrun. Monte Carlo localisation: Efficient position estimation for mobile robots. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 343–349, Orlando, August 1999.

[Frank and Basin, 1998] Ian Frank and David Basin. Search in games with incomplete information: A case study in using Bridge card play. *Artificial Intelligence*, 100(1–2):87–123, 1998.

[Gärdenfors and Makinson, 1988] Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowledge (TARK)*, pages 83–95, Pacific Grove, 1988. Morgan Kaufmann.

[Ginsberg, 2001] Matthew Ginsberg. GIB: Imperfect information in a computationally challenging game. *Journal of Artificial Intelligence Research*, 14:303–358, 2001.

[Grove, 1988] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17(2):157–170, 1988.

[Hollinger and Sukhatme, 2014] Geoffrey Hollinger and Gaurav Sukhatme. Sampling-based robotic information gathering algorithms. *International Journal of Robotics Research*, 39(9):1271–1287, 2014.

[Katsuno and Mendelzon, 1992] Hirofumi Katsuno and Alberto Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1992.

[Lang and Zanuttini, 2012] Jerôme Lang and Bruno Zanuttini. Knowledge-based programs as plans — the complexity of plan verification. In L. De Raedt, C. Bessière, D. Dubois, P. Doherty, P. Frasconi, F. Heintz, and P. Lucas, editors, *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, pages 504–509, Montpellier, France, August 2012. IOS Press.

[McCarthy, 1963] John McCarthy. *Situations and Actions and Causal Laws*. Stanford Artificial Intelligence Project, Memo 2, Stanford University, CA, 1963.

[Reiter, 1991] Ray Reiter. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In V. Lifschitz, editor, *Artificial Intelligence and Mathematical Theory of Computation*, pages 359–380. Academic Press, 1991.

[Reiter, 2001] Ray Reiter. On knowledge-based programming with sensing in the situation calculus. *ACM Transactions on Computational Logic*, 2(4):433–457, 2001.

[Richards and Amir, 2012] Mark Richards and Eyal Amir. Information set generation in partially observable games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Toronto, 2012.

[Scherl and Levesque, 2003] Richard Scherl and Hector Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144(1):1–39, 2003.

[Schofield *et al.*, 2012] Michael Schofield, Timothy Cerexhe, and Michael Thielscher. HyperPlay: A solution to general game playing with imperfect information. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1606–1612, Toronto, July 2012. AAAI Press.

[Shapiro *et al.*, 2000] Steven Shapiro, Maurice Pagnucco, Yves Lespérance, and Hector Levesque. Iterated belief change in the situation calculus. In A. Cohn, F. Giunchiglia, and B. Selman, editors, *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 527–538, Breckenridge, April 2000. Morgan Kaufmann.