Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (2008)

An AGM-Based Belief Revision Mechanism for Probabilistic Spatio-Temporal Logics

Austin Parker, Guillaume Infantes, V.S. Subrahmanian,
austinjp@cs.umd.edu, infantes@umiacs.umd.edu, vs@cs.umd.edu
Department of Computer Science
University of Maryland, College Park

John Grant*
jgrant@towson.edu
Department of Mathematics
Towson University

Abstract

There is now extensive interest in reasoning about moving objects. A PST knowledge base is a set of PST-atoms which are statements of the form "Object o is/was/will be at location L at time t with probability in the interval [L,U]". In this paper, we study mechanisms for belief revision in PST-KBs. We propose multiple methods for revising PST-KBs. These methods involve finding maximally consistent subsets, as well as changing the spatial, temporal, and probabilistic components of the atoms. We show that some methods cannot satisfy the AGM axioms for belief revision, while others do but are coNP-hard. Finally we present an algorithm for revision through probability change which runs in polynomial time and satisfies the AGM axioms.

Introduction

There are numerous applications where we need to reason about probabilistic spatio-temporal applications. A shipping company may be interested in continuously tracking the locations of its vehicles. As RFID tags become ever more common, companies (pharma, automotive, electronics) are interested in tracking supply items and in understanding where these items are now, and where they might be in the future. Military agencies are interested in tracking where vehicles might be - now and in the future. Cell phone companies are interested in when and where cell phones might be in the future in order to determine how best to balance load on cell towers. Moreover, all these applications have an essential component involving uncertainty. Predicting where a cell phone might be in the future may be derived probabilistically from past logs showing the phones' location. Likewise, predicting where and when an RFID tag will be is subject to uncertainty. Where and when a ship will reach a given geolocation is also subject to many forces that cannot be accurately specified, even when a schedule is available.

Methods to reason about probabilistic spatio-temporal (PST) information have emerged in recent years, both in databases (Parker, Subrahmanian, and Grant 2007) and in AI (Cohn and Hazarika 2001; Muller 1998).

One important aspect of applications such as those mentioned above is that there is *continuous change*. As objects move, they encounter unexpected situations, leading to a continuous revision of estimates of where they might be in the future, as well as a revision of where they might have been in the past. Surprisingly, to date, we are not aware of any effort to handle revisions to such PST knowledge bases. A PST knowledge base $\mathcal{K}$ can be revised in many different ways. Clearly, when the insertion of a fact a into the knowledge base leads to no inconsistency, i.e. $\mathcal{K} \cup \{a\}$ is consistent, then a can just be added to $\mathcal{K}$. However, when $\mathcal{K} \cup \{a\}$ is inconsistent, then many different belief revision operations are possible based on whether we modify temporal information, or probabilistic information, or spatial information.

In this paper, we first formalize the problem of *inserting* facts into PST knowledge bases. We also recall the AGM belief revision postulates (Alchourrón, Gärdenfors, and Makinson 1985). We then examine four different ways of revising PST knowledge bases. For each such method, we study which of the AGM postulates it satisfies, as well as what the complexity of the method is. Surprisingly, by revising only the probabilistic aspect of the data, we find that one can satisfy *all* the basic AGM axioms in polynomial time.

Background: Formal Model

(Parker, Subrahmanian, and Grant 2007) proposes a framework for probabilistic spatio-temporal reasoning. However, they place no restrictions on the speed at which a vehicle can travel, nor do they restrict where a vehicle can go. This is clearly unrealistic as no vehicle can travel at arbitrary speeds, and some vehicles cannot go some places (e.g. a car cannot drive across an ocean). PST KBs proposed here enhance their framework by including velocity and reachability constraints. We assume the existence of some set ID of object id's and a finite convex set S of points in a 2-dimensional space[1]. We assume that the set T of time points consists of all non-negative integers. Occasionally, we will make the *bounded time assumption* that T is the set of all non-negative integers upto some arbitrary but fixed maximum time bound.

*John Grant is also affiliated with the Department of Computer Science at the University of Maryland.

[1]The framework is easily extensible to higher dimensions.

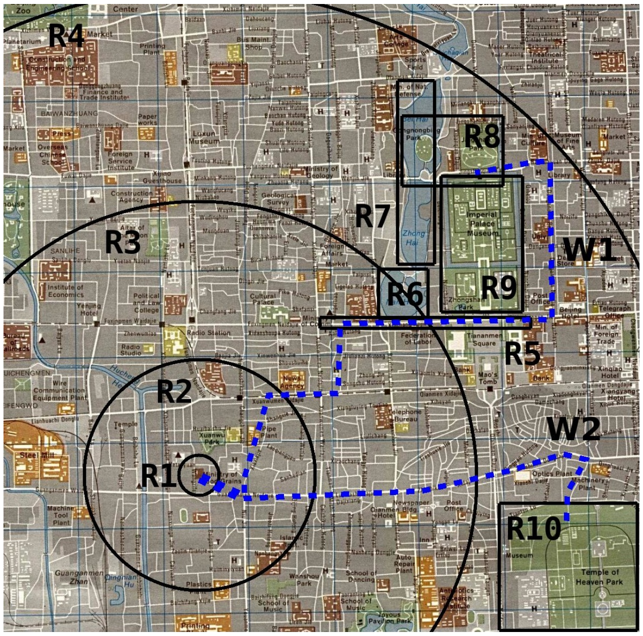

Figure 1: An example PST knowledge base representing possible locations of a delivery boy in Beijing, China. The dotted lines represent possible paths taken by the delivery boy where each dot is the boy's location at a particular time.

Definition 1. *If $id \in ID, t \in T, r \subseteq S$ ($r \neq \emptyset$) and $0 \leq \ell \leq u \leq 1$, then $(id, r, t, [\ell, u])$ is called a **PST-atom**.*

Intuitively, a PST-atom $(id, r, t, [\ell, u])$ says that the object with the given id is somewhere in (or expected to be somewhere in) region r at time t with a probability in the $[\ell, u]$ interval. We use statistical probabilities, make no independence assumptions, and use intervals and linear programs for joint probability computations.

Example 1. *Figure 1 shows regions R1...R10 in Beijing. A pizza shop in the center of R1 is delivering pizzas to the Imperial Palace Museum. The shop guarantees delivery in 40 minutes and wants to reason about the probability of delivering the pizza on time. Delivery boy db leaves at time 0 from region R1, giving the PST-atom: $(db, R1, 0, [1, 1])$. R6 is water-logged, making it impossible for db to be there: so we have the atoms $(db, R6, 0, [0, 0]), \dots, (db, R6, 40, [0, 0])$. We expect the delivery boy to be in region R5 at times 10 to 25 with $45 - 55\%$ probability: $(db, R5, 10, [0.45, 0.55]), \dots, (db, R5, 25, [0.45, 0.55])$. We are almost certain that the delivery boy will not enter the museum, giving the atoms $(db, R9, 1, [0, 0.01]), \dots, (db, R9, 40, [0, 0.01])$. Sometimes db plays hooky and visits the park: $(db, R10, 30, [0, 0.2])$. We use $\mathcal{K}_{Beijing}$ to denote these PST-atoms.*

A *reachability definition RD*, is a set of atoms of the form $reachable(id, p_1, p_2)$ indicating that id can go from point p_1 to point p_2 in one unit of time. Reachability definitions can account for different types of moving objects (e.g. planes vs. bicycles) and different terrain conditions. Moreover, RD does not need to be explicitly stored - it can be implemented through a call to a piece of software code (e.g. Google Maps) that merely returns "true" or "false" when invoked with a triple (id, p_1, p_2). We assume that for every id, the transitive closure of RD is true for every pair of points – no two points are allowed to be completely disconnected.

Example 2. *Given each object's maximal speed v^+_{id}, we define a reachability predicate which requires the object to move at a rate less than v^+_{id}: $reachable(id, p_1, p_2)$ iff $(d(p_1, p_2) < v^+_{id})$.*

Definition 2 (PST-KB)**.** *A **PST-knowledge base** is a pair $(\mathcal{K}, RD)$ where $\mathcal{K}$ is a finite set of PST-atoms and RD is a reachability definition.*

Given a PST-KB $\mathcal{K}$, we use the notation $\mathcal{K}^{id,t}$ to denote the set of all PST atoms of the form $(id, -, t, -)$ in $\mathcal{K}$. *Throughout this paper, we assume the existence of an arbitrary, but fixed reachability definition, RD — hence, we will abuse notation and simply refer to $\mathcal{K}$ as a PST-KB.* We define semantics through worlds.

Definition 3 (World)**.** *A world w is a function, $w : ID \times T \rightarrow S$ such that for all objects id, points p_1, p_2, and time point t if $w(id, t) = p_1$ and $w(id, t+1) = p_2$ then $reachable(id, p_1, p_2) \in RD$. $\mathcal{W}$ is the set of all worlds.*

An interpretation assigns a probability to each world.

Definition 4 (Interpretation)**.** *An interpretation I is a probability distribution over $\mathcal{W}$.*

Intuitively, $I(w)$ is the probability that w describes the actual locations of the objects.

Example 3. *Two paths the delivery boy may take are shown in Fig. 1 as dotted lines W1 and W2. These are potential worlds where the dots give the delivery boy's locations at successive time points. An example interpretation I assigns probability 0.9 to world W1, probability 0.1 to world W2, and probability 0 to any other world.*

The definition of satisfaction of a PST-atom by an interpretation is as follows.

Definition 5 (Satisfaction/Entailment)**.** *Interpretation I satisfies $(id, r, t, [\ell, u])$ (denoted $I \models (id, r, t, [\ell, u])$) iff:*

$$\sum_{w \in \mathcal{W}, w(id,t) \in r} I(w) \in [\ell, u].$$

I satisfies knowledge base $\mathcal{K}$ (denoted $I \models \mathcal{K}$) iff I satisfies all $a \in \mathcal{K}$. $\mathcal{K}$ entails knowledge base $\mathcal{K}'$ (or atom a) iff all I satisfying $\mathcal{K}$ also satisfy $\mathcal{K}'$ (resp. a).

$\mathcal{K}$ is consistent iff there is an interpretation I that satisfies it. $\mathcal{K}$ and $\mathcal{K}'$ are equivalent (denoted $\mathcal{K} \equiv \mathcal{K}'$) iff for all interpretations I, $I \models \mathcal{K}$ iff $I \models \mathcal{K}'$.

A PST-atom atom a *is consistent with* PST-KB $\mathcal{K}$ iff $\mathcal{K} \cup \{a\}$ is consistent.

Example 4. *The atom $(db, R7, 15, [0.75, 0.75])$ is not consistent with the knowledge base $\mathcal{K}_{Beijing}$ from Example 1 because according to $\mathcal{K}_{Beijing}$ db is in region R5 at time 15 with probability in $[0.45, 0.55]$ and R1 is disjoint from R7. The total probability of db on the map at time 15 would then exceed 1.*

However, $(db, R7, 15, [0.45, 0.55])$ *is consistent with* $\mathcal{K}_{Beijing}$ *– consider for instance an interpretation that gives to the delivery boy probability* 0.51 *to be in region R7 and* 0.49 *to be in region R5.*

Consistency Checking

We can check consistency by solving a linear program. Because linear programs can be solved in time polynomial in their input, consistency checking will run in polynomial time when the number of time points is bounded *a priori*.

The linear program we use contains variables of the form $v_{id,t,p,q}$, each representing the probability that object id will be at point p at time t and then at point q at time $t+1$. While this may seem overly complicated, less complicated variable schemes – such as ones where each variable represents the probability that a given object is at a given location at a given time, used in (Parker, Subrahmanian, and Grant 2007) were not easily extendable to handling the intricacies of the reachability predicate.

For convenience, let $minT(\mathcal{K})$ be the minimum time point referenced in $\mathcal{K}$ and $maxT(\mathcal{K})$ be the maximum time point referenced in $\mathcal{K}$. Note that when the bounded time assumption is made, we have *a priori* bounds for $minT(\mathcal{K})$ and $maxT(\mathcal{K})$. We will use an extra timepoint $maxT(\mathcal{K})+1$ for ease of presentation.

Definition 6 (LP($\mathcal{K}$))**.** *We let the linear constraints for* $\mathcal{K}$ *be the set* $LP(\mathcal{K})$ *containing exactly the following constraints for all id referenced in* $\mathcal{K}$*, and all integers t s.t.* $minT(\mathcal{K}) \leq t \leq maxT(\mathcal{K})$*:*

- *For all* $(id, r, t, [\ell, u]) \in \mathcal{K}$*:*
  - $\ell \leq \sum_{p \in r} \sum_{q \in S} v_{id,t,p,q}$ *and* $u \geq \sum_{p \in r} \sum_{q \in S} v_{id,t,p,q}$
- *For all* id, t*:* $\sum_{p \in S} \sum_{q \in S} v_{id,t,p,q} = 1$
- *For all* $p, q \in S$ *and all* id, t*:* $v_{id,t,p,q} \geq 0$
- *For all* $p, q \in S$ *and all* id, t *such that* $\neg reachable(id, p, q)$*:* $v_{id,t,p,q} = 0$.
- *For all* $p \in S$ *and* id, t*:* $\sum_{q \in S} v_{id,t,q,p} = \sum_{q \in S} v_{id,t+1,p,q}$

Theorem 1. $LP(\mathcal{K})$ *has a solution iff* $\mathcal{K}$ *is consistent.*

Proof. Proof sketch:
($\Rightarrow$): Let θ be a solution to $LP(\mathcal{K})$. To construct a satisfying interpretation I, let $\alpha[id, t, p]$ be the probability that id is at p at time t. This can be computed from θ as follows: $\alpha[id, t, p] = \sum_{q \in S} v_{id,t,p,q}\theta$. Define I for all $w \in \mathcal{W}$ s.t. $I(w) = \prod_{t=minT(\mathcal{K})}^{maxT(\mathcal{K})} \alpha[id, t, w(id, t)]$. I is a valid probability distribution over $\mathcal{W}$ because each $\sum_{p \in S} \alpha[id, t, p] = 1$. I also satisfies $\mathcal{K}$: consider for $(id, r, t, [\ell, u]) \in \mathcal{K}$:

$$\sum_{w(id,t) \in r} I(w) = \sum_{p \in r} \alpha[id, t, p] = \sum_{p \in r} \sum_{q \in S} v_{id,t,p,q}$$

Since any solution to $LP(\mathcal{K})$ enforces that $\ell \leq \sum_{p \in r} \sum_{q \in S} v_{id,t,p,q} \leq u$ it follows that $\ell \leq \sum_{w(id,t) \in r} I(w) \leq u$.

($\Leftarrow$): Let I be an interpretation satisfying $\mathcal{K}$. Let θ be an assignment to the variables v such that, for all id, t, p, q:

$$v_{id,t,p,q}\theta = \sum_{w(id,t)=p \wedge w(id,t+1)=q} I(w). \tag{1}$$

θ is also a solution to $LP(\mathcal{K})$. Consider for each $(id, r, t, [\ell, u])$ that $\ell \leq \sum_{w(id,t) \in r} I(w) \leq u$ implies $\ell \leq \sum_{p \in r} \sum_{q \in S} v_{id,t,p,q}\theta \leq u$. That $\sum_{q \in S} v_{id,t+1,p,q}\theta = \sum_{q \in S} v_{id,t,p,q}\theta$ follows from algebraic manipulation. That $\sum_{p \in S} \sum_{q \in S} v_{id,t,p,q} = 1$ for any id, t follows from the fact that $\sum_{w \in \mathcal{W}} I(w) = 1$. That θ solves the rest of the constraints in $LP(\mathcal{K})$ is straightforward. $\square$

The theorem yields a straightforward consistency checking algorithm: simply check if $LP(\mathcal{K})$ has a solution using standard linear programming solvers.

To determine the running time of this algorithm, we count the number of variables and equations in $LP(\mathcal{K})$. The number of variables is dependent upon the number of IDs in the knowledge base, which is at most $|\mathcal{K}|$, the number of points in space, which is constant, and the number of time points $n_t = maxT(\mathcal{K}) - minT(\mathcal{K})$. This gives an upper bound of $O(|\mathcal{K}| \cdot n_t)$ variables. The number of constraints in $LP(\mathcal{K})$ is $2 \times |\mathcal{K}|$ for the constraints from the atoms, plus one constraint per id and t, plus $|S|^2$ constraints per id and T, plus $|S|$ constraints per id and t giving $O((|\mathcal{K}| \cdot n_t)^2)$ constraints (since $|S|$ is constant and since the number of ids is bounded by $|\mathcal{K}|$). Since linear programs are solvable in polynomial time (L.G.Khachiyan 1979), and the input to our linear program solver will be a polynomial in $O((|\mathcal{K}| \cdot n_t)^3)$. n_t is, in general, unbounded. However, if we make the bounded time assumption (for example, assuming a bound of 1000 years may be more than enough for most government and business applications, but not enough for certain applications involving astronomical bodies), then consistency checking is polynomial in the size of the input knowledge base.

Some belief revision strategies

We now present AGM-style postulates (Alchourrón, Gärdenfors, and Makinson 1985) for revising PST-KBs. A revision operator $\dot{+}$ is a binary function that takes a PST-KB and a PST-atom as input, and produces a PST-KB as output. $\dot{+}$ is required to satisfy the AGM axioms[2] expressed in our framework.

(A1) $\mathcal{K} \dot{+} a$ is PST-KB.

(A2) $\mathcal{K} \dot{+} a \models a$.

(A3) $(\mathcal{K} \cup \{a\}) \models (\mathcal{K} \dot{+} a)$.

(A4) If a is consistent with $\mathcal{K}$ then $(\mathcal{K} \dot{+} a) \models (\mathcal{K} \cup \{a\})$.

(A5) $\mathcal{K} \dot{+} a$ is inconsistent iff $\{a\}$ is inconsistent.

(A6) If $a \equiv a'$ then $\mathcal{K} \dot{+} a \equiv \mathcal{K} \dot{+} a'$.

The reader can easily see that the revision of a PST-KB $\mathcal{K}$ with a PST-atom a may be handled in many different

[2]As PST-KBs are atomic, we do not discuss AGM axioms involving negation and disjunction.

ways when $\mathcal{K} \cup \{a\}$ is inconsistent. For example, we could change the t part of a PST-atom, or the r part of a PST-atom, or the $[\ell, u]$ part of a PST-atom.[3] We could also study maximal consistent subsets (Baral, Kraus, and Minker 1991; Fagin, Ullman, and Vardi 1983).

Maximal Consistent Subsets

We can define a revision operator $\dot{+}_m$ based on maximal consistent subsets as follows. For this section, we assume the time points available to be bounded to some fixed, finite, set of integers T.

Definition 7. *Suppose $\mathcal{K}$ is a PST-KB and a is a PST-atom. Then* $\mathcal{K}' \cup \{a\}$ accomplishes the revision of $\mathcal{K}$ by adding a via the subset strategy *iff $\mathcal{K}'$ is a subset of $\mathcal{K}$ and $\mathcal{K}' \cup \{a\}$ is consistent.*

$\mathcal{K}' \cup \{a\}$ optimally accomplishes the revision of $\mathcal{K}$ by adding a via the max-subset strategy *iff it accomplishes the revision of $\mathcal{K}$ by adding a via the subset strategy and there is no other $\mathcal{K}'' \cup \{a\}$ that accomplishes the same revision such that $\mathcal{K}' \subsetneq \mathcal{K}''$.*

We use the notation $\mathcal{K} \dot{+}_m a$ to denote a $\mathcal{K}' \cup \{a\}$ that optimally accomplishes the revision of $\mathcal{K}$ by adding a via the max-subset strategy. [4]

We verify that $\dot{+}_m$ satisfies the AGM axioms.

Proposition 1. *Any function $\dot{+}_m$ that optimally accomplishes the revision via the max-subset strategy satisfies the AGM axioms.*

Unfortunately, computing $\dot{+}_m$ is intractable.

Theorem 2. *Determining if $\mathcal{K}' \cup \{a\}$ optimally accomplishes the revision of $\mathcal{K}$ by adding a via the max-subset strategy is coNP-complete under the bounded time assumption, and coNP-hard otherwise.*

That this problem is in coNP under the bounded time assumption follows from the polynomial time consistency checking algorithm. A witness $\mathcal{K}''$ which is a strict superset of $\mathcal{K}'$ and for whom $\mathcal{K}'' \cup \{a\}$ is consistent proves the given $\mathcal{K}'$ does not optimally accomplish the revision of $\mathcal{K}'$ with a via the max-subset strategy. As consistency checking of PST-KBs is polynomial under the bounded time assumption, this establishes membership in coNP.

To see why it is coNP-hard, consider the coNP-complete case of the knapsack problem defined by $(W = \{w_i\}, c, X = \{x_i\})$ with n items where item j has weight w_j, all items have value 1 and item j is included in the knapsack when $x_j = 1$ and not included when $x_j = 0$. Deciding if $X = \{x_1, \ldots, x_n\}$ maximizes $\sum_{j=1}^n x_j$ subject to $\sum_{j=1}^n w_j x_j \leq c$, with $x_j \in \{0, 1\}$ is coNP-complete. This decision problem reduces to deciding optimal max-subset revision. Construct a max-subset revision problem instance using a space composed of $n+1$ points $\{p_1, ..., p_n, p_{n+1}\}$ a knowledge base $\mathcal{K} = \{(id, \{p_i\}, t, [\frac{w_i}{\sum_i w_i}, \frac{w_i}{\sum_i w_i}]) | 1 \leq i \leq n\}$, a revising atom $a = (id, \{p_i | 1 \leq i \leq n\}, t, [0, \frac{c}{\sum_i w_i}])$, and a revised knowledge base $\mathcal{K}' = \{(id, \{p_i\}, t, [\frac{w_i}{\sum_i w_i}, \frac{w_i}{\sum_i w_i}]) | x_i = 1\}$. X solves the given knapsack problem iff $\mathcal{K}' \cup \{a\}$ optimally accomplishes revision of $\mathcal{K}$ with a via the max-subset strategy.

Minimizing Spatial Change

One may think that we can revise $\mathcal{K}$ by changing the spatial component r of PST atoms in $\mathcal{K}$. A *spatial revision* of PST-atom $a = (id, r, t, [\ell, u])$ is an atom of the form $a' = (id, r', t, [\ell, u])$. The distance $d_S(a, a')$ is given by $abs(|r \cup r'| - |r \cap r'|)$. A spatial revision of PST-KB $\mathcal{K}$ is a knowledge base $\mathcal{K}'$ containing at most one spatial revision of each atom in $\mathcal{K}$. The distance between a PST-KB and its spatial revision ($d_S(\mathcal{K}, \mathcal{K}')$) is the sum of the distances between the individual atoms and their associated spatial revision.

Definition 8. *A spatial revision $\mathcal{K}'$ of $\mathcal{K}$ w.r.t. an inserted PST-atom a is* optimal *iff $\mathcal{K}' \cup \{a\}$ is consistent and there is no other spatial revision $\mathcal{K}''$ of $\mathcal{K}$ w.r.t. a such that $\mathcal{K}'' \cup \{a\}$ is consistent and $d_S(\mathcal{K}, \mathcal{K}'') < d_S(\mathcal{K}, \mathcal{K}')$. We use $\mathcal{K} \dot{+}_s a$ to denote an optimal spatial revision $\mathcal{K}'$. (As in the case of the max-subset strategy, there may be multiple optimal spatial revision strategies, see footnote 4).*

Unfortunately, in general, as the following example shows, there may be cases where no spatial revision satisfies AGM axioms **(A1)** and **(A5)**.

Example 5. *Suppose $ID = \{id\}$ and $S = \{p_1, p_2\}$. Let $\mathcal{K} = \{a_1\}$ where $a_1 = (id, \{p_1\}, 0, [0.5, 0.5])$. Let $a = (id, \{p_1\}, 0, [0, 0])$. By* **(A1)***, $\mathcal{K} \dot{+}_s a$ must be a PST-KB. However, $\mathcal{K} \cup \{a\}$ is inconsistent and $\mathcal{K}$ must be revised. There are 2 possible spatially revised KBs depending on which subset of $\{p_1, p_2\}$ is used as the spatial component of a_1: $(id, \{p_2\}, 0, [0.5, 0.5])$ and $(id, \{p_1, p_2\}, 0, [0.5, 0.5])$. None of these atoms is consistent with a. Hence, Axiom* **(A5)** *is violated.*[5]

Notice that the above example holds both for bounded and unbounded sets of timepoints.

Minimizing Temporal Change

In this section, we study what happens when we revise a PST-KB $\mathcal{K} = \{a_1, \ldots, a_n\}$ by changing $a_i = (id_i, r_i, t_i, [\ell_i, u_i])$ to $a'_i = (id_i, r_i, t'_i, [\ell_i, u_i])$. In other words, the only change allowed in a PST-atom is to modify the time stamp. Given a PST-KB $\mathcal{K}$ of the above form, we call such a revised PST-KB a *temporal variant* of $\mathcal{K}$.

The distance between a temporal variant $\{a'_1, \ldots, a'_n\}$ of a PST-KB $\mathcal{K} = \{a_1, \ldots, a_n\}$, denoted $d_T(\mathcal{K}, \mathcal{K}')$ is given by $\sum_{i=1}^n |t_i - t'_i|$.

$\mathcal{K}'$ is called a *temporally optimal variant* of $\mathcal{K}$ w.r.t. an inserted PST-atom a iff (i) $\mathcal{K}' \cup \{a\}$ is consistent, (ii) $\mathcal{K}'$

[3] Another option is to allow changes to the id part of a PST-atom. We do not study this possibility due to space constraints.

[4] There is some non-determinism in this definition. A strict total ordering O_T can be induced on all $\mathcal{K}'$ satisfying the above definition and the minimal element of the strict total ordering can be picked in order to induce determinism. **Throughout the rest of this paper, we assume such a strict total ordering is available.**

[5] Note that this example does not depend upon how the distance function d_S is defined.

is a temporal variant of $\mathcal{K}$ and (iii) there is no other temporal variant $\mathcal{K}''$ of $\mathcal{K}$ such that $\mathcal{K}'' \cup \{a\}$ is consistent and $d_T(\mathcal{K}, \mathcal{K}'') < d_T(\mathcal{K}, \mathcal{K}')$. As in the case with the previous two revision strategies, there can be multiple temporally optimal variants - see footnote 4. We denote this temporally optimal variant of $\mathcal{K}$ w.r.t. atom a by $\dot{+}_t$.

In this section, we do *not* make the bounded time assumption. Should we make the bounded time assumption, a counter-example similar to example 5 would make AGM-compliant temporal revision impossible in the general case.

Theorem 3. *Suppose $\mathcal{K}$ is a PST-KB and a is an insertion. Checking if $\mathcal{K}'$ is a temporally optimal variant of $\mathcal{K}$ is coNP-hard.*

Proof. For space reasons, we only sketch the proof.

We show a reduction from a special coNP-complete decision version of the knapsack problem specified by $(W = \{w_i\}, c, X = \{x_i\})$ where we are given n items with weights $w_1, \ldots, w_n$ and values of 1. Determining if an assignment $X = \{x_i | 1 \leq i \leq n \wedge x_i \in \{0,1\}\}$ maximizes $\sum_{j=1}^{n} x_j$ subject to $\sum_{j=1}^{n} w_j x_j \leq c$ is coNP-complete. We show a reduction from a problem instance (W, c, X) to an instance $(\mathcal{K}, a, \mathcal{K}')$ of the temporally optimal variant problem. Let $S = \{p_1, \ldots, p_{n+1}\}$, and let there be one object id. We will use two time points: 0 and 1, and the reachability predicate is always true. Let $tot = \sum_{j=1}^{n} w_j$ and let $\mathcal{K} = \left\{ \left(id, \{p_i\}, 1, \left[\frac{w_i}{tot}, \frac{w_i}{tot}\right]\right) \middle| 1 \leq i \leq n \right\}$. The revision atom will be $a = \left(id, \{p_1, \ldots, p_n\}, 1, \left[0, \frac{c}{tot}\right]\right)$. Let $\mathcal{K}' = \{(id, \{p_i\}, x_i, [\frac{w_i}{tot}, \frac{w_i}{tot}]) | 1 \leq i \leq n\}$. Since X maximizes $\sum_{j=1}^{n} w_j x_j$ subject to $\sum_{j=1}^{n} x_j \leq c$ iff $\mathcal{K}' \cup \{a\}$ is a temporally optimal variant of $\mathcal{K}$ via a, this problem is coNP-hard. $\square$

The ***TemporalRevision(**$\mathcal{K}, a$**)*** algorithm works via *unary* temporal variants. $(id, r, t', [\ell, u])$ is a unary temporal variant of $(id, r, t, [\ell, u])$ iff $abs(t - t') = 1$. The algorithm creates a search tree - each node N in the search tree has an $N.KB$ field. The root of the search tree is initialized to $Root.KB = \mathcal{K}$. Every child C of a node N is just like N except that exactly one PST atom in $N.KB$ is replaced by a unary temporal variant. Further, each child knowledge base is required to be further (according to d_T) from $\mathcal{K}$ than its parent. When visiting a node N, the algorithm checks if $N.KB \cup \{a\}$ is consistent. By creating and visiting this tree in breadth first order, we are guaranteed that the first node that satisfies this consistency check is an optimal temporal variant of $\mathcal{K}$ that accomplishes the insertion of a.

Theorem 4. *Algorithm **TemporalRevision** is correct, i.e. **TemporalRevision(**$\mathcal{K}, a$**)** returns a temporally optimal variant of $\mathcal{K}$ that accomplishes the insertion of a as long as a is consistent. It returns "error" iff a is inconsistent. Moreover, **TemporalRevision(**$\mathcal{K}, a$**)** satisfies the AGM axioms.*

Minimizing Probability Change

In this section, we propose a belief revision operator that replaces PST-atoms of the form $(id, r, t, [\ell, u])$ in $\mathcal{K}$ by PST-atoms $(id, r, t, [\ell', u'])$ where $[\ell, u] \subseteq [\ell', u']$. In other words, this belief revision operator expands the probability bounds of PST atoms in $\mathcal{K}$ in order to retain consistency when a is added. Obviously, we want to minimize the expansion of the probability interval $[\ell, u]$ to $[\ell', u']$.

Algorithm 1 ***TemporalRevision(**$\mathcal{K}, a$**)*** Search over potential temporal changes to $\mathcal{K}$.

```
If {a} is inconsistent, return "error".
Get new node Root. Set Root.KB = K;
TODO = [ Root ]. {TODO is an ordered list.}
while True do
  Let nextTODO be an empty list.
  {iterate over TODO in order.}
  for N in TODO do
    if N.KB ∪ {a} is consistent return N.KB ∪ {a}.
    Insert each child of N into nextTODO.
  end for
  Let TODO=nextTODO.
  sort TODO with strict total ordering O_T (see footnote 4).
end while
```

Definition 9. *Suppose $a = (id, r, t, [\ell, u])$ is a PST-atom and $[\ell, u] \subseteq [\ell', u']$. Then the PST-atom $a' = (id, r, t, [\ell', u'])$ is called a* weakening *of a. The distance, $d_P(a, a')$ between a and a' is defined as $(\ell - \ell') + (u' - u)$.*

A PST-KB $\mathcal{K}'$ is called a weakening *of a PST-KB $\mathcal{K}$ iff there is a bijection β from $\mathcal{K}$ to $\mathcal{K}'$ such that for all $a \in \mathcal{K}$, $\beta(a)$ is a weakening of a. The distance $d_P(\mathcal{K}, \mathcal{K}')$ between $\mathcal{K}$ and $\mathcal{K}'$ is defined as $\Sigma_{a \in \mathcal{K}} d_P(a, \beta(a))$.*

In most cases, β can be derived directly by manipulating the probability bounds associated with a PST-atom $a \in \mathcal{K}$. In the sequel we assume β is known.

Definition 10. *Suppose $\mathcal{K}$ is a PST-KB and a is a PST-atom. A weakening $\mathcal{K}'$ of $\mathcal{K}$ is called an* optimal weakening *of $\mathcal{K}$ w.r.t. the insertion of a iff: (i) $\mathcal{K}' \cup \{a\}$ is consistent and (ii) for every other weakening $\mathcal{K}''$ of $\mathcal{K}$ such that $\mathcal{K}'' \cup \{a\}$ is consistent, $d_P(\mathcal{K}, \mathcal{K}') \leq d_P(\mathcal{K}, \mathcal{K}'')$.*

We can find an optimal weakening of PST-KBs by setting up a linear program with variables $v_{id,t,p,q}$ each representing the probability of an object id being at location p at time t and at location q at time $t+1$. We limit the range of id to those objects mentioned in the database and the range of t to the bounded set T provided *a priori* (we assume a bounded set of timepoints T for probabilistic revision). For each PST-atom $a_i = (id_i, r_i, t_i, [\ell_i, u_i])$ in $\mathcal{K}$, we also include variables low_i and up_i for the atoms' modified lower and upper bounds.

Definition 11 (Probability Revision Linear Program (PRLP))**.** *Let $PRLP(\mathcal{K}, a)$ contain only the following:*

1. *For each $a_i = (id_i, r_i, t_i, [\ell_i, u_i]) \in \mathcal{K}$:*
   (a) $0 \leq \left(\sum_{p \in r_i} \sum_{q \in S} v_{id_i, t_i, p, q}\right) - low_i$
   (b) $0 \geq \left(\sum_{p \in r_i} \sum_{q \in S} v_{id_i, t_i, p, q}\right) - up_i$.
   (c) $\ell_i \geq low_i$*,* $low_i \geq 0$*,* $u_i \leq up_i$*, and* $up_i \leq 1$
2. *For $a = (id', r', t', [\ell, u])$:*
   (a) $\ell \leq \sum_{p \in r'} \sum_{q \in S} v_{id', t', p, q}$ *and* $u \geq \sum_{p \in r'} \sum_{q \in S} v_{id', t', p, q}$

3. *For each id in the knowledgebase and each t in T*
   (a) *For all $p, q \in S$, $v_{id,t,p,q} \geq 0$.*
   (b) $\sum_{p \in S} \sum_{q \in S} v_{id,t,p,q} = 1$
   (c) *For all $p, q \in S$, if $\neg reachable(id, p, q)$: $v_{id,t,p,q} = 0$*
   (d) *For all $p \in S$: $\sum_{q \in S} v_{id,t,q,p} = \sum_{q \in S} v_{id,t+1,p,q}$*

We now compute an optimal weakening of $\mathcal{K}$ by minimizing the distance function $\Sigma_{a_i \in \mathcal{K}} d_P(a_i, \beta(a_i))$ subject to $PRLP(\mathcal{K}, a)$. As in the case of all our revision strategies, when there are multiple solutions to this linear program, we assume there is a mechanism to deterministically pick one. We are now able to define a probabilistic revision strategy.

Definition 12 (Probabilistic Revision)**.** *Suppose $\mathcal{K}$ is a PST-KB and a is a PST atom. Let θ be a (deterministically) selected solution of the linear program **minimize** $\Sigma_{a_i \in \mathcal{K}}((\ell_i - low_i) + (up_i - u_i))$ **subject to** $PRLP(\mathcal{K}, a)$. Return the PST-KB, denoted $\mathcal{K} \dot{+}_p a$ defined as*

$$\{ (id_i, r_i, t_i, [low_i\theta, up_i\theta]) | (id_i, r_i, t_i, [\ell_i, u_i]) \in \mathcal{K}\} \cup \{sa\} .$$

Since the number of points in space and times points is constant, only the number of objects mentioned in $\mathcal{K}$ and the number of atoms in $\mathcal{K}$ affect the number of variables in $PRLP()$, which is $O(|\mathcal{K}|)$. The number of constraints is similarly limited by $O(|\mathcal{K}|)$. Thus the size of the entire linear program created by PRLP is polynomial in $|\mathcal{K}|$. Since solving linear programs is also polynomial (L.G.Khachiyan 1979), and we can assume our mechanism for picking a solution deterministically runs in polynomial time[6]; hence the above procedure computes $\mathcal{K} \dot{+} a$ in polynomial time.

This polynomial time probabilistic revision strategy also satisfies the requisite AGM axioms.

Proposition 2. *$\mathcal{K} \dot{+}_p a$ satisfies* **(A1)-(A6)**.

Related Work and Conclusion

There is much work on *spatio-temporal logics* (Gabelaia et al. 2003; Merz, Wirsing, and Zappe 2003) in the literature. These logics extend temporal logics to handle space. There is also much work on qualitative spatio-temporal theories (for a survey see (Cohn and Hazarika 2001; Muller 1998)). (Shanahan 1995) discusses the frame problem when constructing a logic-based calculus for reasoning about the movement of objects in a real-valued co-ordinate system. (Rajagopalan and Kuipers 1994) focuses on relative position and orientation of objects with existing methods for qualitative reasoning in a Newtonian framework.

In contrast to these works, we focus on the problem of belief revision in spatio-temporal logics with uncertainty (not handled in past work). We first build on the framework of (Parker, Subrahmanian, and Grant 2007) in order to include the realistic requirement that vehicles have movement constraints and velocity constraints and we show how to handle consistency checking in this setting. We then develop analogs of the AGM axioms to handle insertions into PST-KBs and evaluate different ways of accomplishing these revisions. We show that the max-consistent subset and temporal revision strategies satisfy the AGM axioms, but respectively lead to coNP-complete and coNP-hard problems. Spatial revisions do not satisfy the AGM axioms. Our final result shows that probabilistic revision satisfies the AGM axioms and is polynomially computable making it (to our mind at least), the preferred option.

Future work will need to focus on how to incorporate insertions into PST KBs *efficiently*. This is complex because in applications involving GPS sensors, updates occur continuously, leading to a large volume of updates. Taming this complexity will be quite a challenge and will perhaps need methods that are even more efficient than the polynomial strategy of minimizing probability change.

[6]Such mechanisms clearly exist: consider a strict total ordering over the variables which tells the order with which the linear program solver should minimize variables.

Acknowledgements. Research funded in part by AFOSR grants FA95500610405 and FA95500510298, ARO grant DAAD 190310202 and NSF grant 0540216. We thank Francesco Parisi for assistance with Theorems 2 and 3 and the AAAI reviewers for their helpful comments.

References

Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.

Baral, C.; Kraus, S.; and Minker, J. 1991. Combining multiple knowledge bases. *IEEE TKDE* 3(2):208–220.

Cohn, A. G., and Hazarika, S. M. 2001. Qualitative spatial representation and reasoning: an overview. *Fundam. Inf.* 46(1-2):1–29.

Fagin, R.; Ullman, J. D.; and Vardi, M. Y. 1983. On the semantics of updates in databases. In *PODS*, 352–365. ACM.

Gabelaia, D.; Kontchakov, R.; Kurucz, A.; Wolter, F.; and Zakharyaschev, M. 2003. On the computational complexity of spatio-temporal logics. In *FLAIRS Conference*, 460–464.

L.G.Khachiyan. 1979. A polynomial algorithm in linear programming. *Soviet Mathematics Doklady* 20:191–194.

Merz, S.; Wirsing, M.; and Zappe, J. 2003. A spatio-temporal logic for the specification and refinement of mobile systems. In Pezzè, M., ed., *FASE*, volume 2621 of *Lecture Notes in Computer Science*, 87–101. Springer.

Muller, P. 1998. Space-Time as a Primitive for Space and Motion . In *FOIS*, 63–76. Amsterdam: IOS Press.

Parker, A. J.; Subrahmanian, V.; and Grant, J. 2007. A logical formulation of probabilistic spatial databases. *IEEE TKDE* 19(11):1541–1556.

Rajagopalan, R., and Kuipers, B. 1994. Qualitative spatial reasoning about objects in motion: Application to physics problem solving. In *IJCAI'94*, 238–245.

Shanahan, M. 1995. Default reasoning about spatial occupancy. *Artif. Intell.* 74(1):147–163.