

A Framework for Reasoning Under Uncertainty with Temporal Constraints

Eugene Santos, Jr., Deqing Li, and John T. Wilkinson

Thayer School of Engineering

Dartmouth College

Hanover, N.H., U.S.A.

{Eugene.Santos.Jr, Deqing.Li, John.T.Wilkinson}@Dartmouth.edu

Abstract—Time is the key stimulus to change, causality and interaction which are the main components of a dynamic world. Therefore, the modeling of knowledge, especially in complex and dynamic domains like economics, sociology, and ecology, must incorporate the concept of time. Although there has been much research over the years on the representation of knowledge (causality, implication, and uncertainty) and on the representation of time, it has been a continuing challenge to unify them in a meaningful and useful fashion. In this paper, we propose a framework for reasoning under uncertainty with temporal constraints. The framework is extended from Bayesian knowledge-bases (BKBs), which represent knowledge in an “if-then” structure and represent uncertainty based on probability theory. By adding temporal constraints to BKBs, the framework provides a comprehensive model that incorporates the semantics of both time and uncertainty.

Keywords—temporal reasoning, probabilistic reasoning, knowledge representation, Bayesian knowledge-base.

I. INTRODUCTION

Traditional knowledge representations model knowledge as a set of concepts and the connections between them. However, in domains like economics, sociology, and ecology, entities interact with each other according to not only internal and external stimuli but also the passage of time. Time is ubiquitous and essential in models that seek to determine the roots of change, causality, and interaction. Thus, it ultimately must be a critical component of any knowledge representation. For example, say we have the following knowledge representing a person’s chance of spreading the flu:

If a person who has the flu virus is an adult, they are able to infect others from the time symptoms develop up until the 5th day of showing symptoms. On the other hand, if they are a child, they are contagious as long as they are sick. Additionally, it is sometimes possible for people to spread the flu virus before showing symptoms or after symptoms improve. Symptoms tend to develop before the 4th day of catching the flu and to improve 7 days after the virus gets into the body.

This example depicts the relationship between age and spreading the flu, and the uncertainty in this relationship. Thus, an ideal framework must be able to describe the temporal constraints on the knowledge and the uncertainty present, as

well as reason through all this. By providing observations and evidence from the past or present, the framework should be able to help predict the future state of the world as well as the time that this state may be attained. In addition, it should be able to explain the causes of the event and when they occurred.

Early representations such as first order logic presented significant challenges toward effectively representing time and/or uncertainty. Techniques such as fuzzy logic and Bayesian graphical models were developed to capture uncertainty, while solutions such as Interval Algebra were created to allow the representation of temporal information. Bayesian networks (BNs) [1], which allow representing and reasoning over uncertain knowledge based on probability theory, continue to gain popularity. However, they can require the complete specification of large conditional probability tables and do not permit knowledge cycles. In contrast, Bayesian knowledge-bases (BKBs) [2] are robust to knowledge incompleteness (essentially only the probabilistic knowledge available is included) and cyclic information, while still preserving a probabilistically sound representation of uncertainty. Unfortunately, neither BNs nor BKBs inherently capture temporal information or temporal relationships among pieces of knowledge. While there exist some approaches that attempt to capture probabilistic uncertainty and time (e.g., dynamic BNs and probabilistic temporal networks (PTNs) [3]), they suffer from a variety of limitations in representational power.¹

In this paper, we present a unified framework called temporal Bayesian knowledge-bases (TBKBs) for reasoning under uncertainty with temporal information that is based on BKBs. In this framework, temporal constraints are captured as disjunctive linear systems (DLSs) [4] that allow representation of and reasoning over intervals of time in a flexible manner. Inferencing in TBKBs is founded on probability theory while guided by temporal constraints. In this way, the semantics of interactions and time are managed simultaneously. In order to deal with incompleteness in time, both temporal interactions and atemporal interactions, which are akin to those found in standard BNs and BKBs and not constrained by time, are allowed during reasoning.

¹ PTNs currently come closest to a unified model of time and uncertainty but are ultimately limited by the underlying BN model.

This paper is organized as follows: We will first introduce some background information including theories of temporal reasoning and BKBs. Section 3 details the formal definition of TBKB and its properties, followed by its inference scheme and fusion algorithm. Then to demonstrate the utility of TBKBs, an example illustrating how to construct and reason over a TBKB representing the flu scenario mentioned earlier is provided. Finally, we conclude and discuss future work.

II. BACKGROUND

A. Temporal Reasoning

Time is fundamental to reasoning about change and action in the real world. Our experiences occur over time and our perception of the world uses time as a reference [5]. Researchers took notice of the necessity of modeling change as far back as the 1960's when the earliest study in the AI community can be traced back to McCarthy and Hayes' *situation calculus* [6], but it was not until the 1970's that people began to represent time explicitly. Two of the most influential theories about time are Allen's *Interval Algebra (IA)* [7] and McDermott's *temporal logic* [8]. For this paper, we utilize and extend the former. In IA, each event is associated with a time interval, and any pair of time intervals can be related in one of thirteen ways which are *equals*, *precedes*, *meets*, *during*, *starts*, *finishes*, *overlaps*, and their inverses. Most of the research on temporal knowledge representation [9], [10], [11] is extended from IA with parallel work [12] built on point algebra (PA). We also use intervals for representing time in our approach. A BKB captures possible states of the world and their likelihood. These states may only be attainable if certain temporal relations are satisfied. The qualitative relations of IA are quantitatively represented in TBKBs to facilitate inferencing. For example, the *precedes* relationship between a variable a and a variable b in IA can be represented as the inequality $t_a^- < t_a^+ < t_b^- < t_b^+$ given that the temporal interval of variable x is denoted by (t_x^-, t_x^+) . However, TBKBs are not limited to the thirteen relationships defined in IA. For instance, we could specify that variable b can only be true more than two time units after variable a ceases to be true with the inequality $t_a^+ + 2 < t_b^-$; a relationship that is not representable in IA. Thus, our framework subsumes and extends the representational power of IA. In addition, reasoning about points, as in PA, is possible in our approach. The point t_0 can be represented with the interval $(t_0, t_0 + \epsilon)$, where ϵ is arbitrarily close, but not equal, to zero and is held constant throughout the model [3].

Due to the important role that time plays in the interaction of objects, many domains such as medical diagnosis, planning and scheduling, natural language processing, and object detection and tracking have attempted to incorporate temporal dimensions in their knowledge representation systems [13], [14], [15]. The main components for modeling these domains include temporal reasoning, abductive reasoning, and reasoning under uncertainty. Much recent work in temporal modeling has focused on temporal reasoning with a coarse knowledge of time (also called fuzzy intervals). Examples include using the idea of a neighborhood to constrain coarse time relations [16], handling controllability and preference of temporal constraints simultaneously [17], and using possibility theory to represent uncertainty in temporal relations [12]. However, the semantics

for handling uncertainty in the knowledge and in its interactions were largely ignored. A few of the approaches that take all three components into consideration include Kjaerulff [18], Santos and Young [3], and Hanks et al. [19]. In Kjaerulff [18], the relationship between the previous and the current states of entities are encoded in arcs. Thus, the changing states of the world are dynamically modeled as time-sliced BNs. Santos and Young [3] incorporated temporal intervals and temporal relations directly into BNs, formulating the model as a constraint satisfaction problem. In Hanks et al. [19], events are divided into endogenous (internal) events and exogenous (external) events. Endogenous events influence exogenous ones, and the world is updated according to the time when the next change is triggered.

The framework for representing temporal knowledge proposed in this paper addresses all three components we have identified in modeling a dynamic world. The framework is built upon BKBs which support uncertain causal relationships, knowledge incompleteness and cyclic information. It also allows uncertainty in temporal relationships between knowledge concepts.

B. Bayesian Knowledge Bases

Like a BN, a BKB represents the causal relationships between knowledge using a directed graph composed of nodes and arcs [2]. However, BKBs introduce a more compact representation while still retaining the expressiveness of BNs by allowing partial probability distributions. Nodes in BKBs express knowledge about the world in the form of random variables (r.v.s) and the uncertainty of knowledge in the form of conditional probability rules (CPRs) with edges denoting the causal relationships between the nodes connected by the CPRs. BKBs feature two types of nodes. Instantiation nodes or "I-nodes," depicted as white ovals in Figure 1, express r.v. instantiations (or states), e.g. $A=a$, where A is a r.v. and a is one of its possible states. This is in contrast to a BN where each node represents a r.v. along with all of its instantiations. Instead of links between r.v.s, BKBs use support nodes, or "S-nodes," depicted as black dots in Figure 1, to encode the conditional dependencies between their descendant I-nodes (also called their head) and their predecessor I-nodes (also called their tail). The weight attached to an S-node represents the conditional probability of its head given that its tail is known to be true. Below we present a formal definition of a BKB from Santos and Santos [2].

Definition. A *correlation-graph* is a directed graph $G=(I \cup S, E)$ in which $I \cap S = \Phi$, $E \subset \{I \times S\} \cup \{S \times I\}$, and $\forall a \in S$, there exists a unique $b \in I$ such that $a \rightarrow b \in E$. Φ denotes the empty set, I denotes the set of all I-nodes in graph G , S denotes the set of all S-nodes, E denotes the set of all edges, and $a \rightarrow b$ denotes a directed edge connecting node a to node b . For an S-node a in graph G , we use $Tail_G(a)$ to represent a 's tail, and $Head_G(a)$ to represent a 's head. Two sets of I-nodes, I_1 and I_2 are said to be *mutually exclusive* if there is an I-node $v=i_1$ in I_1 and an I-node $v=i_2$ in I_2 such that $i_1 \neq i_2$.

Definition. A Bayesian knowledge-base (BKB) is a tuple $B = (G, w)$ where $G = (I \cup S, E)$ is a correlation-graph, and $w: S \rightarrow [0,1]$ such that

1. $\forall a \in S$, $Tail_G(a)$ contains at most one instantiation of each r.v.
2. For distinct S-nodes $a, b \in S$ that support the same I-node, $Tail_G(a)$ and $Tail_G(b)$ are mutually exclusive.
3. $\forall Q \subseteq S$ s.t. (i) $Head_G(a)$ and $Head_G(b)$ are mutually exclusive, (ii) $Tail_G(a)$ and $Tail_G(b)$ are not mutually exclusive for all a and b in Q , $\sum_{q \in Q} w(q) \leq 1$.




Figure 1. Example of a BKB fragment.

The explicit representation of rules between states of r.v.s introduced by BKBs is more flexible than in BNs. In a BN, the full conditional probability table must be provided at each node, while in BKBs, only the rules that are available need be specified. Another advantage arises from the BKBs' ability to deal with cyclic knowledge as discussed in [2].

Probabilistic inferencing over a BKB can answer questions like "What is the most probable state of the world", a process referred to as belief revision. In belief revision, the goal is to find the most-probable states of all the r.v.s by maximizing $P(A_{k+1}=a_{k+1}, \dots, A_n=a_n | A_1=a_1, \dots, A_k=a_k)$ where $A_1=a_1, \dots, A_k=a_k$ are evidence. In conducting belief revision, we make use of a special type of subgraph of a BKB called an *inference*. An inference can be formally defined as follows **Error! Reference source not found.** [1]:

Definition. Let $B = (G, w)$ be a BKB with correlation-graph $G = (I \cup S, E)$. A subgraph $\tau = (I' \cup S', E')$ of G , where $I' \subseteq I$, $S' \subseteq S$ and $E' \subseteq E$, is called an *inference* over B if

1. τ is acyclic.
2. (Well-supported) $\forall a \in I', \exists b \in S'$ s.t. $b \rightarrow a \in E'$.
3. (Well-founded) $\forall b \in S', Tail_{\tau}(b) = Tail_G(b)$.
4. (Well-defined) $\forall b \in S', Head_{\tau}(b) = Head_G(b)$.
5. There is at most one I-node corresponding to any given random variable in I' .

The probability of an inference is simply the product of the weights of all S-nodes in the inference. Algorithms for performing belief revision in BKBs are discussed in [20]. Probabilistic inferencing can be applied to prediction and explanation. Prediction is achieved by extending the current evidence forward to the unknown portion of the world, while explanation is to determine the cause of the current observations by extending the evidence backwards to hypotheses.

A. Formulation

BKBs have been applied to a wide variety of domains such as adversary intent inferencing [21], social network analysis [22], and medical information retrieval [23], to name a few. TBKBs provide a finer granularity of relationships, and thus can contribute to more accurate prediction and explanation. A preliminary version of the TBKB framework was described in Pioch et al. [24], where it was used to model a commander's perception of adversaries and other entities in a dynamic battlespace. TBKBs can be simply understood as BKBs that also allow for temporal constraints and temporal evidence. We first define TBKBs and related terminologies, followed by a description of how to inference over a TBKB and how to fuse several smaller TBKBs into one larger TBKB. Finally, we formally derive some of the properties of TBKBs.

Definition. A *temporal interval* is a pair (t, t^+) where t represents the start time, t^+ represents the end time, $t^- < t^+$, and $t^-, t^+ \in R$.

Temporal intervals are used to determine the time during which an I-node may be true. Constraints on these intervals are represented by a DLS.

Definition. A *disjunctive linear system (DLS)* L is a set $\{A_1, A_2, \dots, A_n\}$ where each A_i is an individual system of linear equations and inequalities. $V(L)$ represents the variable set for L such that a variable v is in $V(L)$ iff it has a nonzero coefficient in some linear system A_i in L . A *solution* to L is a complete assignment of values to $V(L)$ that satisfies some A_i in L .

In TBKBs, the constraints are defined on the S-nodes, in terms of the I-nodes adjacent to a given S-node.

Definition. A *temporal causal relationship (TCR)* r_a describes a disjunctive set of temporal relationships between a set of I-nodes X that are adjacent to a single S-node a . r_a is a temporal causal relationship iff r_a can be expressed as a DLS of time intervals of I-nodes in X .

For example, suppose there are n I-nodes in X , whose temporal intervals are denoted as (t_x^-, t_x^+) , then r_a is a TCR on the nodes in X iff r_a can be expressed as $\{A_1, A_2, \dots, A_n\}$ where

$$A_i = \begin{cases} \alpha_{i,1,1}t_1^- + \alpha_{i,1,2}t_1^+ + \dots + \alpha_{i,1,2n-1}t_n^- + \alpha_{i,1,2n}t_n^+ \Delta c_{i,1} \\ \dots \\ \alpha_{i,j,1}t_1^- + \alpha_{i,j,2}t_1^+ + \dots + \alpha_{i,j,2n-1}t_n^- + \alpha_{i,j,2n}t_n^+ \Delta c_{i,j} \\ \dots \end{cases}$$

where $\alpha_{i,j,2k-1}$ denotes the coefficient of t_k^- , $\alpha_{i,j,2k}$ denotes the coefficient of t_k^+ , $c_{i,j}$ denotes the constant in the j^{th} linear inequality, and Δ is one of the following relations: $<$, $>$, \geq , $=$ or \neq .

The representation of a TCR as a DLS is highly flexible and expressive. The thirteen popularly used relations in IA, such as the *meets* relation, can be fully represented, and extended, by DLSs. For example, the temporal relationship "(I-node x_l meets I-node y and the duration of y is as long as the sum of the

duration of x_1 and that of x_2) or (y starts at least 3 hours after x_1 finishes)” can be expressed as the following DLS.

$$r = \left\{ \begin{array}{l} t_{x_1}^+ = t_y^- \\ t_y^+ - t_y^- = (t_{x_1}^+ - t_{x_1}^-) + (t_{x_2}^+ - t_{x_2}^-) \\ t_y^- > t_{x_1}^+ + 0300 \end{array} \right\} \vee$$

Note that this duration constraint would not be allowed in IA. r is not strictly in the format described above, but can be easily rewritten to conform to that format.

Definition. A *temporal Bayesian knowledge-base (TBKB)* is a pair $T = (B, R)$ where $B = (G, w)$ is a Bayesian knowledge-base with $G = (I \cup S, E)$, and R is a set of temporal causal relationships defined on S , i.e. $\forall r_a \in R, a \in S$.

B. An Example

In this section, we build a TBKB modeling the flu spread scenario described earlier in order to demonstrate the construction and inferencing techniques of TBKBs. The underlying BKB represents the causal relationships between elements that contribute to the flu spread as shown in Figure 2. It says, for example, that if an adult has developed symptoms of the flu, there is an 83% chance that he is contagious. Temporal dependencies are shown on the S-nodes. For example, when the symptoms develop will depend on the time that the person catches the flu. According to the scenario, symptoms tend to develop within 4 days of the flu virus entering the body, and do not improve until 7 days after catching the flu. It is necessary to encode the temporal relationships both between getting the flu and showing symptoms and between getting the flu and not showing symptoms, the first dependency can be represented by $\{t_{Symp=Y}^- \leq t_{Flu=Y}^- + 4\}$, and the second one by $\{t_{Symp=N}^+ \leq t_{Flu=Y}^- + 4\} \vee \{t_{Symp=N}^- \geq t_{Flu=Y}^- + 7\}$. The second relationship uses a DLS to represent the “or” condition because “not showing symptoms” happens both at the beginning and at the end of getting the flu. We further notice that there are some implicit dependencies in the scenario, namely that symptoms appear only after the virus enters the body and disappear before the person has fully recovered. As a result of combining all the dependencies, we end up with the TBKB shown in Figure 2. You may notice that temporal relationships are absent on some S-nodes, which is due to the fact that temporal dependencies are not applicable such as between I-nodes $Age=Adult$ and $Flu=Yes$, or that information is not available such as between variables $Symptoms=No$ and $Contagious=Yes$, but we have an implicit relationship that $Contagious=Yes$ happens no earlier than $Symptoms=No$.

IV. TEMPORAL BKB REASONING

Since TBKBs are built upon BKBs, techniques for reasoning and fusion in BKBs can be extended to accommodate TBKBs as well. Below, we present algorithms for reasoning and fusion in TBKBs.

Notation: Let Σ_v denote the set of all states of random variable v .

Definition. a is a *temporal assignment (TA)* iff a is a triple (v, γ, ε) where v is a random variable, γ is a temporal interval, and $\varepsilon \in \Sigma_v$.

One of the key components that is needed to extend BKB reasoning to the temporal domain is a definition of an inference in a TBKB.

Definition. Let $T = (B, R)$ be a temporal Bayesian knowledge-base. A *temporal inference* over T is a pair $\tau = (\tau', \alpha)$ where $\tau' = (I' \cup S', E')$ is an inference over B ,

$$\alpha = \left\{ \alpha_i = (v_i, \gamma_i, x_i) \mid \bigcup_i x_i = I' \text{ and } \bigcap_i x_i = \Phi \right\}$$

is a set of temporal assignments that is a solution to all DLSs representing the TCRs on S' .

That is to say, a temporal inference is an inference with a temporal assignment for each I-node in the inference such that the set of temporal assignments is a solution to all DLSs present on the S-nodes in the inference.

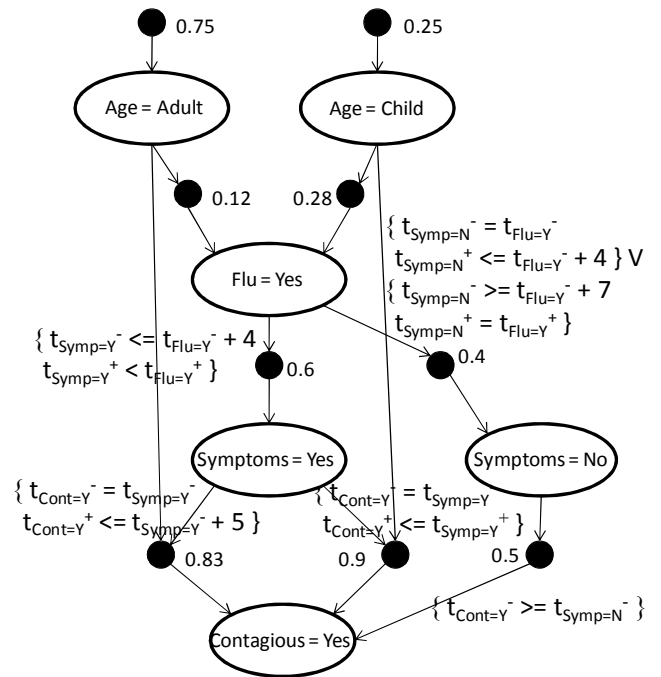


Figure 2. a TBKB example showing causes of spreading flu

As in the case of an atemporal inference, the probability of a temporal inference is the product of the weights of all S-nodes in the inference. The definition above ensures that the temporal inference is valid in terms of both probability and time. In a traditional BKB, evidence takes the form of the assignment of random variables to states. In TBKBs several types of evidence are available. This includes setting a random variable to a state, setting the starting point or endpoint of a temporal interval for an I-node, and introducing new temporal constraints between multiple I-nodes. This allows for random

variables to be assigned to states during specific times and to further constrain the solution space.

We now describe how to perform belief revision on a TBKB. The idea behind the algorithm is to find the temporal inference that has maximum probability among the set of all possible temporal inferences containing the temporal evidence. To accomplish this, we perform the belief revision algorithm that is used on atemporal BKBs, but at each step of the algorithm, we ensure that the temporal constraints are satisfiable and we in fact produce a set of temporal assignments that satisfy the constraints if so. The temporal assignments are obtained by solving a *cluster constraint system (CCS)* [4] for each inference τ .

Definition. A *cluster constraint system* C is a set $\{L_1, L_2, \dots, L_n\}$ where each L_i is a DLS. A *solution* to C is an assignment of the variables in C to values such that the assignment is a solution to each L_i . Thus C is a conjunction of DLSs.

In this case, we form a CCS from the set of all TCRs in τ . We also add two additional linear systems to the CCS respectively guaranteeing that the start of each interval is before its end and that the start of any interval in the tail of an S-node is before the start of the interval of the S-node's head (causality). Finally, we add the temporal evidence in as another DLS. If there exists a solution to the CCS, the solution serves as a set of TAs that satisfies all temporal constraints in the scenario and contains all provided evidence. Although solving a CCS is NP-hard, interval reasoning problems using CCSs can be efficiently approximated by the *convex envelope* technique proposed in [4].

We now present an example of belief revision based on the TBKB in Figure 2. We want to ask what the state of a person is given that they caught the flu on day 0 and spread the flu to others on day 10. Assuming that it is well known that the duration of flu symptoms is 7 days, two pieces of evidence ($Flu=Yes$ and $Contagious=Yes$) and three pieces of temporal evidence ($t_{Flu=Y}^- = 0$, $t_{Sym=Y}^+ - t_{Sym=Y}^- = 7$ and $t_{Con=Y}^- \leq 10 \leq t_{Con=Y}^+$) are available. The most probable inference of the underlying BKB is that $Age=Adult$, $Flu=Yes$, $Symptoms=Yes$, and $Contagious=Yes$ with joint probability 0.045. However, the temporal constraints between an adult showing flu symptoms and being contagious indicate that the person can be contagious until day 9, which contradicts the evidence of spreading the flu on day 10. Therefore, this is not a valid temporal inference and so it is eliminated. The second most probable underlying BKB inference is $Age=Child$, $Flu=Yes$, $Symptoms=Yes$, and $Contagious = Yes$ which has a valid temporal assignment for each node at $(-\infty, \infty)$, $(0, \infty)$, $(4, 11)$, $(4, 11)$ respectively. By the usage of the symbol ∞ , we mean that any number that forms a valid temporal interval will satisfy the system. Moreover, valid temporal assignments may not be unique, e.g. $Symptoms = Yes$ is possible at both $(3, 10)$ and $(4, 11)$. In the case of the I-node $Flu=yes$, there are infinitely many valid assignments. We arbitrarily choose one of them currently, but may improve the scheme in the future.

Extended from BKBs, TBKBs inherit all the advantages of BKBs such as compact representation, easy construction, and the ability to handle incomplete and cyclic information.

Support for incompleteness is also present in the representation of temporal knowledge in TBKBs. In some cases, complete temporal information is not attainable, e.g. the light was observed to be turned off at 9 o'clock, but when it was turned on is unknown. In this case, the assignment of partial temporal evidence is allowed. Another type of incompleteness is a lack of temporal relationships, e.g., two hours after Jim came into the lab, the light was turned off, but no information about how long before Jim came that the light was turned on can be obtained. Thus, only the temporal relations that are available will be used in reasoning.

V. TEMPORAL BKB FUSION

In Santos et al. [25], an algorithm was proposed to “fuse” several atemporal BKBs into a single larger BKB. In the context of the fusion algorithm, we refer to each input BKB as a *fragment* since it represents one modular piece of knowledge. Fusing fragments while retaining probabilistic correctness is not as straightforward as it may seem at first, especially if the fragments were created in isolation, or by different knowledge engineers. For example, the engineers may disagree on the direction of causality between two nodes or the appropriate magnitude of the conditional probability on an S-node, among other things. To solve this problem, we can parameterize the fragments with additional variables that are referred to as *source nodes*, indicating the source and reliability of the fragment [25].

We note that an analogous algorithm also works on TBKBs. In this context we refer to the fragments as *temporal fragments*. To fuse them we take the union of the input fragments by merging I-nodes that are identical, but keeping all S-nodes, even if they represent the same rules. If an S-node whose head is a state of r.v. A came from fragment σ , we add an I-node $S_A = \sigma$ to its tail along with an S-node supporting the new I-node whose weight is the reliability of the fragment σ . An example of fusing TBKBs is shown in Figure 3. In the example, TBKBs from two sources 1 and 2 are fused into one larger TBKB shown on the right. For each S-node in the original TBKBs, a source node and a reliability node are attached. We assign equal reliability to the sources in this case. The common I-node $A=a$ is merged in the final TBKB, but all S-nodes and temporal relationships remain separated.

VI. PROPERTIES OF TBKBs

We now present some provable properties of TBKBs. First, we are interested in the probabilistic soundness of TBKBs.

Theorem 1. For any set of mutually incompatible temporal inferences² R' in a TBKB T , $\sum_{\tau \in R'} P(\tau) \leq 1$.

Sketch of Proof. We note that the probability of a temporal inference is equal to the probability of its underlying atemporal inference. According to Santos and Santos [2, Key Theorem 4.3], for any set of mutually incompatible inferences R in a BKB B , $\sum_{\tau \in R} P(\tau) \leq 1$. Thus any set of mutually incompatible

² Incompatible inferences have I-node sets that are mutually exclusive.

temporal inferences can be shown to satisfy this inequality by considering the set of underlying atemporal inferences. ■

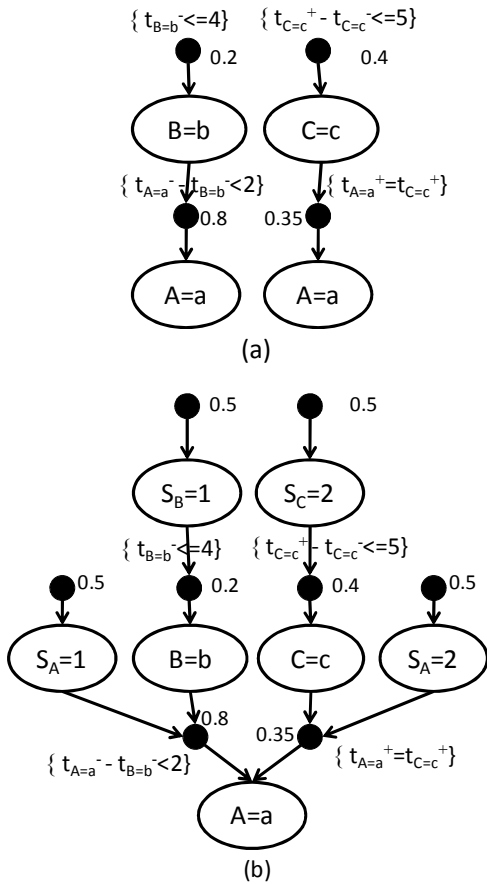


Figure 3. Example of TBKF fusion. (a) Original BKB fragments. (b) Fused BKB fragment.

Another important property of BKBs is groundedness. If an S-node in a BKB remains grounded and un-altered, then even if the rest of the BKB changes, the initial probability assigned to it is semantically preserved [26]. We want to investigate whether groundedness is preserved as information is added to a TBKB as well.

Definition. A node $a \in I \cup S$ in a TBKB is said to be *grounded* if there exists a temporal inference τ over the TBKB such that a is in τ . A TBKB T is said to be *grounded* if $\forall a \in I \cup S, a$ is grounded.

Proposition 1. If all S-nodes in a TBKB T are grounded, then all I-nodes in it are grounded.

To see that this is true, let a be an I-node in a TBKB. Then a must either be in the head or in the tail of some S-node b , and since b is grounded, a must be grounded as well since the head and tail of each S-node in an inference must also be in the inference.

Proposition 2. If a TBKB T is grounded, then its underlying BKB B is grounded, but the reverse is not true.

The first half of *Proposition 2* is obvious. The second half is true because combining the TCRs from each inference into

one CCS may result in a system that is not satisfiable making a valid atemporal inference an invalid temporal inference.

Theorem 2. In TBKB fusion, if the input TBKBs are grounded, then the fused TBKB T' is grounded.

Sketch of Proof. For any S-node $a \in S_i$ from the input TBKB T_i , there is a temporal inference τ over T_i such that $a \in \tau$. We will construct a new subgraph $\tau' \supset \tau$ in T' , and claim that τ' is a temporal inference over T' such that $a \in \tau'$.

To construct τ' , for any S-node $b \in \tau$ whose head is $A=a$, we include its source node $S_A = i$ together with the source node's supporting S-node in τ' . Firstly, τ' is acyclic since τ is acyclic and added nodes (source nodes and supports) don't introduce cycles. Since no new I-nodes are added without a supporting S-node, τ' is well-supported. τ' is also well-founded because the S-nodes added to τ , which are the supports of the source nodes, have empty tails. Moreover, τ' is well-defined since S-nodes are added to support their respective source nodes. Finally, all added I-nodes are source nodes that are not contained in the original TBKBs, and their supporting S-nodes do not introduce new temporal constraints. Therefore, τ' is a temporal inference. ■

VII. CONCLUSION AND FUTURE WORK

This paper proposed a unified framework for reasoning under uncertainty with time by incorporating temporal information into Bayesian knowledge-bases. Its reasoning scheme and fusion algorithm were naturally extended from those used on traditional BKBs. Inferencing over temporal Bayesian knowledge-bases remains probabilistically sound and preserves the semantics in representing uncertainty in the interaction between variables and uncertainty in the time these interactions may occur. Fusion of TBKBs was shown to produce a valid TBKB which remains grounded if the original TBKBs are grounded. As a whole, TBKBs allow the modeling of higher dimensional knowledge by providing a sound representation of time-constraints, which is more appropriate than atemporal knowledge representations in modeling real-world scenarios.

Our future work will focus on studying the limits of the current framework and improving its flexibility in modeling various domains. For example, we will allow assignment of multiple intervals to a r.v. to address cases where an event may be true in discontinuous chunks of time. Moreover, it is possible to allow contradicting states of a r.v. to be true as long as the temporal assignments of the states are mutually exclusive. This indicates that the r.v. switches from one state to another over time. In addition, we will further investigate the properties of TBKBs, in particular, regarding the preservation of groundedness through insertion and modification of temporal knowledge and fusion of fragments.

ACKNOWLEDGMENT

This work was supported in part by a grant from the Air Force Research Laboratory, Information Directorate, Office of Naval Research Grant No. N00014-06-C-0020, and Air Force Office of Scientific Research Grant Nos. FA9550-06-1-0169 and FA9550-07-1-0050.

REFERENCES

- [1] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*. San Francisco, CA: Morgan Kaufmann Publishers, 1988.
- [2] E. Santos, Jr. and E. S. Santos, "A framework for building knowledge-based under uncertainty", in *Journal of Experimental and Theoretical Artificial Intelligence*, vol. 11, pp. 265-286, 1999.
- [3] E. Santos, Jr. and J. D. Young, "Probabilistic temporal networks: a unified framework for reasoning with time and uncertainty", in *International Journal of Approximate Reasoning*, vol. 20, pp. 263-291, 1999.
- [4] E. Santos, Jr., E. E. Santos and K. Kim, "Satisfying constraint sets through convex envelopes", in *Journal of Experimental & Theoretical Artificial Intelligence*, vol. 18, no. 3, pp. 413-432, Sep. 2006.
- [5] L. Vila, "A survey on temporal reasoning in Artificial Intelligence", in *AI Communications*, vol. 7, issue 1, pp. 4-28, Mar. 1994.
- [6] J. McCarthy and P. J. Hayes, "Some philosophical problems from the standpoint of Artificial Intelligence", in *Machine Intelligence 4*, pp. 463-502, 1969.
- [7] J. F. Allen, "Maintaining knowledge about temporal intervals", in *Communications of the ACM*, vol. 26, issue 11, pp. 832-843, 1983.
- [8] D. McDermott, "A temporal logic for reasoning about processes and plans", in *Cognitive Science*, vol. 6, pp. 101-155, 1982.
- [9] G. Ligozat, "Weak Representation of Interval Algebra", in *Proc. AAAI Conference*, pp. 715-720, 1990.
- [10] M. Vilain and H. Kautz, "Constraint propagation algorithms for temporal reasoning", in *Proc. AAAI Conference*, pp. 377-382, 1986.
- [11] S. Schockaert and M. D. Cock, "Temporal reasoning about fuzzy intervals", in *Artificial Intelligence*, vol. 172, issue 8-9, pp. 1158-1193, May 2008.
- [12] D. Dubois, A. Hadjali and H. Prade, "A possibility theory-based approach for handling of uncertain relations between temporal points", in *Proc. 11th International Symposium on Temporal Representation and Reasoning*, pp. 36-43, 2004.
- [13] D. Moldovan, C. Clark and S. Harabagiu, "Temporal context representation and reasoning", In *Proc. Nineteenth International Joint Conference on Artificial Intelligence*, 2005.
- [14] Chilukuri K. Mohan, Kishan G. Mehrotra, Pramod K. Varshney and Jie Yang, "Temporal uncertainty reasoning networks for evidence fusion with applications to object detection and tracking", in *Information Fusion*, vol. 8, issue 3, pp. 281-294, July 2007.
- [15] J. F. Allen, "Planning as temporal reasoning", in *Proc. KR'91*, pp. 3-14, 1991.
- [16] C. Freska, "Temporal reasoning based on semi-intervals", in *Artificial Intelligence*, vol. 54, pp. 199-227, 1992.
- [17] F. Rossi, K. B. Venable and N. Yorke-Smith, "Uncertainty in soft temporal constraint problems: a general framework and controllability algorithms for the fuzzy case", in *Journal of AI Research*. v. 27, pp. 617-674, 2006.
- [18] U. Kjaerulff, "Dhugin: a computational system for dynamic time-sliced Bayesian networks", in *International Journal of Forecasting*, vol. 11, pp. 89-111, 1995.
- [19] S. Hanks, D. Madigan and J. Gavrín, "Probabilistic temporal reasoning with endogenous change", in *Proc. Eleventh Annual Conference on Uncertainty in Artificial Intelligence (UAI-95)*, Aug. 1995.
- [20] Rosen, Tzachi, Shimony, Solomon Eyal, and Santos, Eugene, Jr., "Reasoning with BKBs – Algorithms and Complexity," *Annals of Mathematics and Artificial Intelligence* 40(3-4), 403-425, 2004.
- [21] E. Santos, Jr., Q. Zhao, F. Pratto, A. R. Pearson, B. McQueary, A. Breeden, and L. S. Krause, "Modeling Multiple Communities of Interest for Interactive Simulation and Gaming: The Dynamic Adversarial Gaming Algorithm Project", In *Proc. SPIE: Defense & Security Symposium*, vol. 6564, Orlando, FL, 2007.
- [22] E. E. Santos, E. Santos, Jr., L. Pan and J. T. Wilkinson, "Culturally-infused social network analysis," in *Proc. International Conference on Artificial Intelligence*, 2008.
- [23] E. Santos, Jr., H. Nguyen, Q. Zhao and E. Pukinskis, "Empirical Evaluation of Adaptive User Modeling in a Medical Information Retrieval Application", in *Lecture Notes in Artificial Intelligence*, vol. 2702, User Modeling 2003, P. Brusilovsky, A. Corbett, and F. de Rosis, Eds. Berlin/ Heidelberg: Springer, pp. 292-296, 2003.
- [24] N. J. Pioch, J. Melhuish, A. Seidel, E. Santos, Jr., D. Li and M. Gorniak, "Adversarial intent modeling using embedded simulation and temporal Bayesian knowledge bases", *Proc. SPIE Defense, Security and Sensing 2009*, vol. 7348, Orlando, FL, 2009.
- [25] E. Santos, Jr., J. T. Wilkinson and E. E. Santos, "Bayesian knowledge fusion", *Proc. 22nd International FLAIRS Conference*, pp. 559-564, Sanibel Island, FL, May 2009.
- [26] E. Santos, Jr., E. S. Santos, S. E. Shimony, "Implicitly preserving semantics during incremental knowledge base acquisition under uncertainty", in *International Journal of Approximate Reasoning*, vol. 33, issue 1, pp. 71-94, Apr. 2003.