

Specific-to-General Learning for Temporal Events with Application to Learning Event Definitions from Video

Alan Fern
Robert Givan
Jeffrey Siskind

AFERN@PURDUE.EDU
GIVAN@PURDUE.EDU
QOBI@PURDUE.EDU

School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA

Abstract

We develop, analyze, and evaluate a novel specific-to-general learner for a simple temporal logic, using the resulting algorithm to learn visual event definitions from video sequences. First, we introduce a simple, propositional, temporal, event-description language called AMA that is sufficiently expressive to represent many events yet sufficiently restrictive to support learning. We then give algorithms along with lower and upper complexity bounds for the subsumption and generalization problems for AMA formulas. We present a positive-examples-only specific-to-general learning method based on these algorithms. We also present a polynomial-time-computable “syntactic” subsumption test that implies semantic subsumption without being equivalent to it. A generalization algorithm based on syntactic subsumption can be used in place of semantic generalization to improve the asymptotic complexity of the resulting learning algorithm. Finally, we apply this algorithm to the task of learning event definitions from video and show that it yields definitions that are competitive with hand-coded ones.

1. Introduction

Humans conceptualize the world in terms of objects and events. This is reflected in the fact that we talk about the world using nouns and verbs. We perceive events taking place between objects, we interact with the world by performing events on objects, and we reason about the effects that actual and hypothetical events performed by us and others have on objects. We also *learn* new object and event types from novel experience. In this paper, we present and evaluate novel implemented techniques that allow a computer to learn new event types from examples. We show results from an application of these techniques to learning new event types from automatically constructed relational, force-dynamic descriptions of video sequences.

We wish the acquired knowledge of event types to support multiple modalities. Humans can observe someone *faxing* a letter for the first time and quickly be able to recognize future occurrences of faxing, perform faxing, and reason about faxing. It thus appears likely that humans use and learn event representations that are sufficiently general to support fast and efficient use in multiple modalities. A long-term goal of our research is to allow similar cross-modal learning and use of event representations. We intend the same learned representations to be used for vision (as described in this paper), planning (something that we are beginning to investigate), and robotics (something left to the future).

A crucial requirement for event representations is that they capture the *invariants* of an event type. Humans classify both picking up a cup off a table and picking up a dumbbell off the floor as *picking up*. This suggests that human event representations are *relational*. We have an abstract relational notion of *picking up* that is parameterized by the participant objects rather than distinct propositional notions instantiated for specific objects. Humans also classify an event as *picking up* no matter whether the hand is moving slowly or quickly, horizontally or vertically, leftward or rightward, or along a straight path or circuitous one. It appears that it is not the characteristics of participant-object motion that distinguish *picking up* from other event types. Rather, it is the fact that the object being picked up changes from being supported by resting on its initial location to be supported by being grasped by the agent. This suggests that the primitive relations used to build event representations are *force dynamic* (Talmy, 1988).

Another desirable property of event representations is that they be *perspicuous*. Humans can introspect and describe the defining characteristics of event types. Such introspection is what allows us to create dic-

tionaries. To support such introspection, we prefer a representation language that allows such characteristics to be explicitly manifest in event definitions and not emergent consequences of distributed parameters as in neural networks or hidden Markov models.

We develop a learner for an event representation possessing these desired characteristics as follows. First, we present a simple, propositional, temporal logic called AMA that is a sublanguage of a variety of familiar temporal languages (e.g. linear temporal logic, or LTL (Bacchus & Kabanza, 2000), temporal event logic (Siskind, 2001)). This logic is expressive enough to describe a variety of interesting temporal events, but restrictive enough to support an effective learner, as we demonstrate below. We proceed to develop a specific-to-general learner for the AMA logic by giving algorithms and complexity bounds for the subsumption and generalization problems involving AMA formulas. While we show that semantic subsumption is intractable, we provide a weaker syntactic notion of subsumption that implies semantic subsumption but can be checked in polynomial time. Our implemented learner is based upon this syntactic subsumption.

We next show means to adapt this (propositional) AMA learner to learn relational concepts. We evaluate the resulting relational learner in a complete system for learning force-dynamic event definitions from positive-only training examples given as real video sequences. This is not the first system to perform visual event recognition from video. We review prior work and compare it to the current work later in the paper. In fact, two such prior systems have been built. HOWARD (Siskind & Morris, 1996) learns to classify events from video using temporal, relational representations. But these representations are not force dynamic. LEONARD (Siskind, 2001) classifies events from video using temporal, relational, force-dynamic representations but does not learn these representations. It uses a library of hand-code representations. This work adds a learning component to LEONARD, essentially duplicating the performance of the hand-coded definitions automatically.

While we have demonstrated the utility of our learner in the visual event learning domain, we note that there are many domains where interesting concepts take the form of structured temporal sequences of events. In machine planning, macro-actions represent useful temporal patterns of action. In computer security, typical application behavior, represented perhaps as temporal patterns of system calls, must be differentiated from compromised application behavior (and likewise authorized user behavior from intrusive behavior).

In what follows, Section 2 introduces our application domain of recognizing visual events. Section 3 describes the high-level construction of our learner. Section 4 introduces the AMA language, syntax and semantics, and several concepts needed in our analysis of the language. Section 5 develops and analyzes algorithms for the subsumption and generalization problems in the language, and introduces the more practical notion of “syntactic subsumption”. Section 6 extends the basic propositional learner to handle relational data, negation, and to control exponential run-time growth. Section 7 presents our results on visual event learning, and Sections 8 and 9 compare to related work and conclude.

2. Recognizing Visual Events

LEONARD (Siskind, 2001) is a system for recognizing visual events from video camera input—an example of a simple visual event is “a hand picking up a block”. This research was originally motivated by the problem of adding a learning component to LEONARD. Below we briefly describe the system and the framework for extending LEONARD to learn to recognize events.

LEONARD is a three-stage pipeline depicted in Figure 1. The raw input consists of a video-frame sequence depicting events. First, a segmentation-and-tracking component transforms this input into a polygon movie: a sequence of frames, each frame being a set of convex polygons placed around the tracked objects in the video. Figure 2a shows a “partial” video sequence of a *PICKUP* event that is overlaid by the corresponding polygon movie. Next, a model-reconstruction component transforms the polygon movie into a force-dynamic model. This model describes the changing support, contact, and attachment relations between the tracked objects over time. Figure 2b shows a visual depiction of the force-dynamic model corresponding to the *PICKUP* event. Finally, an event-recognition component armed with a library of event definitions determines which events occurred in the model and, accordingly, in the video. Figure 2c shows the “text” output and input of the event-recognizer for the *PICKUP* event. The first line corresponds to the output which indicates the




Figure 1: The upper boxes represent the three primary components of LEONARD’s pipeline. The lower box depicts the event-learning component described in this paper. The input to the learning component consists of training models of a target event (e.g., movies of *PICKUP* events) and the output is an event definition (e.g., a temporal logic formula defining *PICKUP*).

interval(s) where a *PICKUP* occurred—the remaining lines are the text encoding of the event-recognizer input (model-reconstruction output), indicating the time-intervals in which various force-dynamic relations are true in the video.

The event-recognition component of LEONARD represents event types with event logic formulas like the following simplified example, representing *x* picking up *y* off of *z*.

$$\text{PICKUP}(x, y, z) \triangleq (\text{SUPPORTS}(z, y) \wedge \text{CONTACTS}(z, y)); (\text{SUPPORTS}(x, y) \wedge \text{ATTACHED}(x, y))$$

This formula asserts that an event of *x* picking up *y* off of *z* is defined as a sequence of two states where *z* supports *y* by way of contact in the first state and *x* supports *y* by way of attachment in the second state. SUPPORTS, CONTACTS, and ATTACHED are primitive force-dynamic relations. This formula is a specific example of the more general class of AMA formulas that we use in our learning, presented later in Section 4.

Prior to the work reported in this paper, the definitions in LEONARD’s event recognition library were hand coded. Here, we add a learning component to LEONARD so that it can learn to recognize events. Figure 1 shows how the learning component fits into the overall system. The input to the learning component consists of force-dynamic models from the model-reconstruction stage and its output consists of event definitions which are used by the event-recognizer. We take a supervised learning approach where the force-dynamic model-reconstruction process is applied to training videos of a target event-type—the resulting force-dynamic models are then given to the learner which induces a candidate definition of the event-type. Note that our learning component does not require negative examples of the event type (i.e., movies depicting non-occurrences).

3. Bottom-up Learning from Positive Data

In this work we present a specific-to-general positive-only learner for temporal events—our learning algorithm is only given positive training examples (where the target event occurs) and is not given negative examples (where the target event does not occur). The positive-only setting is of interest as it appears that humans are able to learn many event definitions given primarily or only positive examples. From a practical standpoint, a positive-only learner removes the often difficult task of collecting negative examples that are “representative” of what is not the event to be learned.

A typical learning domain specifies an example space (the objects we wish to classify) and a concept language (formulas that represent sets of examples that they *cover*). Generally we say a concept C_1 is more general (less specific) than C_2 if and only if C_2 is a subset of C_1 —alternatively, a generality relation that may not be equivalent to subset may be specified, often for computational reasons. Setting the goal of finding a concept consistent with a set of positive-only training data generally results in the trivial solution of returning

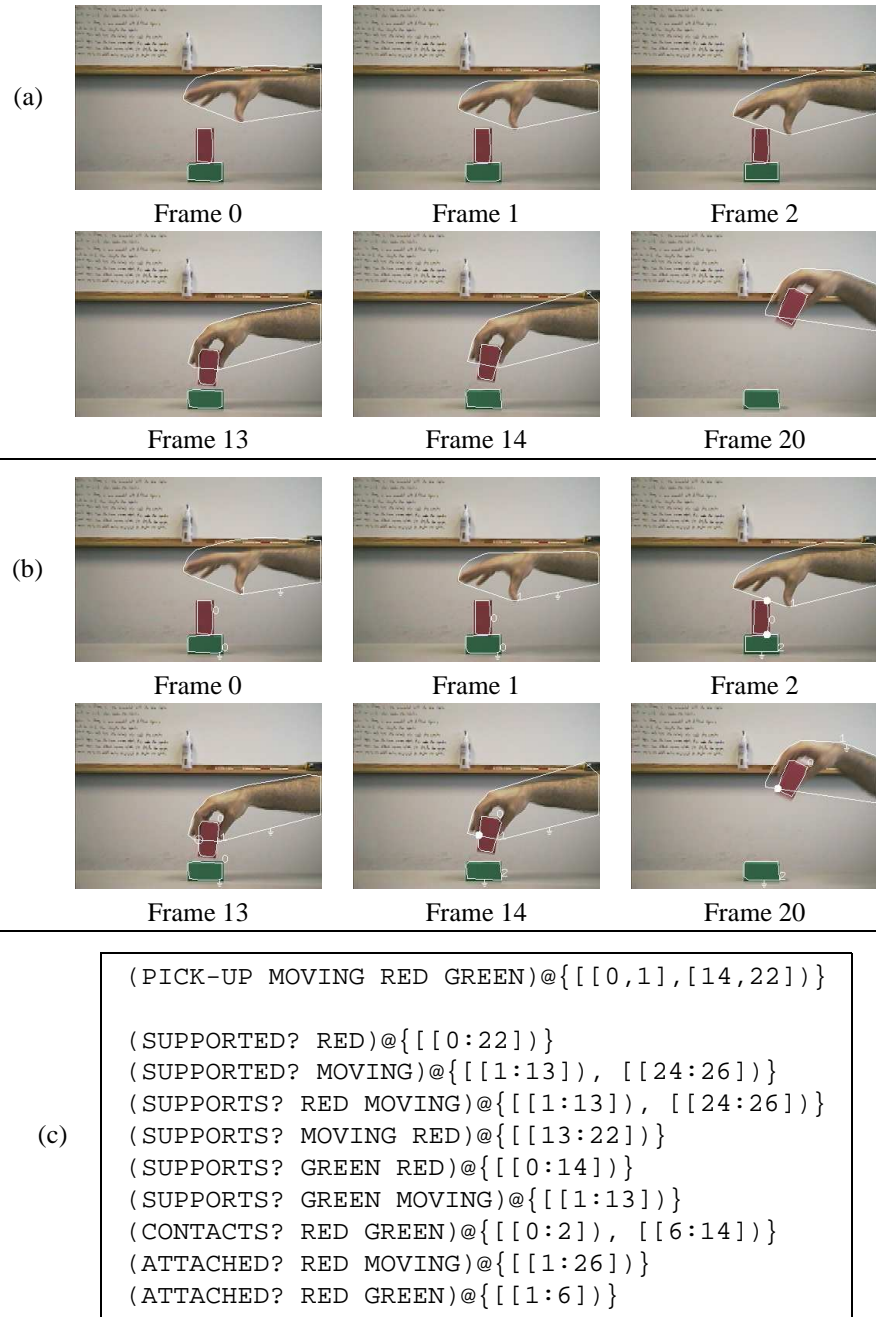


Figure 2: LEONARD recognizes a *PICKUP* event. (a) Frames from the raw video input with the automatically generated polygon movie overlaid. (b) The same frames with a visual depiction of the automatically generated force-dynamic properties. (c) The text input and output of the event classifier corresponding to the depicted movie. The top line is the output and the remaining lines make up the input that encodes the changing force-dynamic properties over time.

the most general concept in the language. To avoid adding negative training data, it is common to specify the learning goal as finding the least-general concept that covers all of the data¹. With enough data and an appropriate concept language, the least-general concept often converges usefully.

We take a standard specific-to-general machine-learning approach to finding the least-general concept covering a set of positive examples. Assume we have a concept language L and an example space S . The approach relies on the computation of two functions: the least-general covering formula (LGCF) of an example and the least-general generalization (LGG) of a set of formulas. An LGCF in L of an example in S is a formula in L that covers the example such that no other covering formula is strictly less general. Intuitively, the LGCF of an example, if unique, is the “most representative” formula in L of that example. An LGG of any subset of L is a formula more general than each formula in the subset and not strictly more general than any other such formula. Neither the LGG nor the LGCF is guaranteed to exist or be unique—these properties must be shown for any language of interest.

Given the existence and uniqueness (up to concept equivalence) of the LGCF and LGG, the specific-to-general approach proceeds as follows. First, use the LGCF to transform each positive training instance into a formula of L . Second, return the LGG of the resulting formulas. The returned formula represents the least-general concept in L that covers all the positive training examples. This learning approach has been pursued for a variety of concept languages including, clausal first-order logic (Plotkin, 1971), definite clauses (Muggleton & Feng, 1992), and description logic (Cohen & Hirsh, 1994). It is important to choose an appropriate concept language as a bias for this learning approach or the concept returned may simply be (or resemble) one of two extremes, either the disjunction of the training data or the universal concept.

In this work, the concept language is the AMA temporal event logic presented below and the example space is the set of all models of that logic. Intuitively, a training example depicts a model where a target event occurs. (The models can be thought of as movies.) We will consider two notions of generalization for AMA concepts (semantic generalization and a weaker syntactic counterpart) and, under both notions, study the properties and computation of the LGCF and LGG.

4. Representing Events with AMA

We study a subset of an interval-based logic called *event logic* (Siskind, 2001) utilized by LEONARD for event recognition in video sequences. This logic is “interval-based” in explicitly representing each of the possible interval relationships given originally by Allen (1983) in his calculus of interval relations (e.g., “overlaps”, “meets”, “during”). Event logic formulas allow the definition of event-types which can specify static properties of intervals directly and dynamic properties by hierarchically relating sub-intervals using the Allen relations. In this paper the formal syntax and semantics of full event logic are needed only for Proposition 4 and are given in Appendix A.

Here we restrict our attention to a much simpler subset of event logic we call AMA, defined below. We believe that our choice of event logic rather than first-order logic, as well as our restriction to the AMA fragment of event logic, provide a useful learning bias by ruling out a large number of ‘practically useless’ concepts while maintaining substantial expressive power. The practical utility of this bias is demonstrated via our empirical results in the visual event recognition application. AMA can also be seen as a restriction of LTL to conjunction and “Until”, with similar motivations. Below we present the syntax and semantics of AMA along with some of the key technical properties of AMA that will be used throughout this paper.

4.1 AMA Syntax and Semantics

It is natural to describe temporal events by specifying a sequence of properties that must hold over consecutive time-intervals; e.g., “a hand picking up a block” might become “the block is not supported by the hand and then the block is supported by the hand.” We represent such sequences with *MA timelines*², which are

1. In some cases, there can be more than one such least-general concept. The set of all such concepts is called the “specific boundary of the version space” (Mitchell, 1982).
 2. MA stands for “Meets/And”, an MA timeline being the “Meet” of a sequence of conjunctively restricted intervals.

sequences of conjunctive state restrictions. Intuitively, an MA timeline is given by a sequence of propositional conjunctions, separated by semi-colons, and is taken to represent the set of events that temporally match the sequence of consecutive conjunctions. An AMA formula is then the conjunction of a number of MA timelines, representing events that can be simultaneously viewed as satisfying each of the conjoined timelines. Formally, the syntax of AMA formulas is given by,

$$\begin{aligned}
 \textit{state} & ::= \textit{true} \mid \textit{prop} \mid \textit{prop} \wedge \textit{state} \\
 \textit{MA} & ::= (\textit{state}) \mid (\textit{state}); \textit{MA} \quad // \textit{may omit parens} \\
 \textit{AMA} & ::= \textit{MA} \mid \textit{MA} \wedge \textit{AMA}
 \end{aligned}$$

where *prop* is any primitive proposition (sometimes called a primitive event-type). We take this grammar to formally define the terms “MA timeline”, “MA formula”, “AMA formula”, and “state”. A *k*-MA formula is an MA formula with at most *k* states, and a *k*-AMA formula is an AMA formula all of whose MA timelines are *k*-MA timelines. We often treat states as proposition sets (with **true** the empty set) and AMA formulas as MA timeline sets. We may also treat MA formulas as sets of states—it is important to note, however, that MA formulas may contain duplicate states, and the duplication can be significant. For this reason, when treating MA timelines as sets, we formally intend sets of *state-index pairs* (where the index gives a states position in the formula). We do not indicate this explicitly to avoid encumbering our notation, but the implicit index must be remembered whenever handling duplicate states.

The semantics of AMA formulas is defined in terms of temporal models. A temporal model $\mathcal{M} = \langle \mathcal{M}, \mathcal{I} \rangle$ over the set of propositions PROP is a pair of a mapping *M* from the natural numbers (representing time) to the truth assignments over PROP, and a closed natural number interval *I*. We note that Siskind (2001) gives a continuous-time semantics for event logic where the models are defined in terms of real-valued time-intervals. The temporal models defined here use discrete natural-number time-indices; however, our results here still apply under the continuous-time semantics (that semantics bounds the number of state changes in the continuous timeline to a countable number). It is important to note that the natural numbers in the domain of *M* are representing time discretely, but that there is no prescribed unit of continuous time represented by each natural number. Instead, each number represents an arbitrarily long period of continuous time during which nothing changed. Similarly, the “states” in our MA timelines represent arbitrarily long periods of time during which the conjunctive restriction given by the state holds. The satisfiability relation for AMA formulas is given as follows:

- A state *s* is satisfied by model $\langle M, I \rangle$ iff $M[x]$ assigns *P* true for every $x \in I$ and $P \in s$.
- An MA timeline $s_1; s_2; \dots; s_n$ is satisfied by a model $\langle M, [t, t'] \rangle$ iff there exists some t'' in $[t, t']$ such that $\langle M, [t, t''] \rangle$ satisfies s_1 and either $\langle M, [t'', t'] \rangle$ or $\langle M, [t'' + 1, t'] \rangle$ satisfies $s_2; \dots; s_n$.
- An AMA formula $\Phi_1 \wedge \Phi_2 \wedge \dots \wedge \Phi_n$ is satisfied by \mathcal{M} iff each Φ_i is satisfied by \mathcal{M} .

The condition defining satisfaction for MA timelines may appear unintuitive at first due to the fact that there are two ways that $s_2; \dots; s_n$ can be satisfied. The reason for this becomes clear by recalling that we are using the natural numbers to represent continuous time intervals. Intuitively, from a continuous-time perspective, an MA timeline is satisfied if there are consecutive continuous-time intervals satisfying the sequence of consecutive states of the MA timeline. The transition between consecutive states s_i and s_{i+1} can occur either within an interval of constant truth assignment (that happens to satisfy both states) or exactly at the boundary of two time intervals of constant truth value. In the above definition, these cases correspond to $s_2; \dots; s_n$ being satisfied during the time intervals $[t'', t']$ and $[t'' + 1, t']$, respectively.

When \mathcal{M} satisfies Φ we say *M* is a model of Φ . We say AMA Ψ_1 *subsumes* AMA Ψ_2 iff every model of Ψ_2 is a model of Ψ_1 , written $\Psi_2 \leq \Psi_1$, and we say Ψ_1 *properly subsumes* Ψ_2 , written $\Psi_2 < \Psi_1$, when we also have $\Psi_1 \not\leq \Psi_2$. Alternatively, we may state $\Psi_2 \leq \Psi_1$ by saying that Ψ_1 is *more general (or less specific) than* Ψ_2 or that Ψ_1 *covers* Ψ_2 . Siskind (2001) provides a method to determine whether a given model satisfies a given AMA formula.

Finally, it will be useful to associate a distinguished MA timeline to a model. The *MA-projection* of a model $\mathcal{M} = \langle M, [i, j] \rangle$ (written as $\text{MAP}(\mathcal{M})$) is an MA timeline $s_0; s_1; \dots; s_{j-i}$ where state s_k gives the true propositions in $M(i+k)$ for $0 \leq k \leq j-i$. Later we show that the MA-projection of a model can be viewed as “representing” that model in a precise sense.

The following two examples illustrate some basic behaviors of AMA formulas.

Example 1 (Stretchability) *The MA timelines $S_1; S_2; S_3$, $S_1; S_2; S_2; \dots; S_2; S_3$, and $S_1; S_1; S_2; S_3; S_3; S_3$ are all equivalent. In general, MA timelines have the property that duplicating any state results in a formula equivalent to the original formula. Recall that, given a model $\langle M, I \rangle$, we view each truth assignment $M[x]$ as representing a continuous time-interval—this interval can conceptually be divided into an arbitrary number of subintervals. Thus if state S is satisfied by $\langle M, [x, x] \rangle$, then so is the state sequence $S; S; \dots; S$.*

Example 2 (Infinite Descending Chains) *Given propositions A and B , the MA timeline $\Phi = (A \wedge B)$ is subsumed by each of the formulas $A; B$, $A; B; A; B$, $A; B; A; B; A; B$, $A; B; A; B; A; B; A; B$, \dots . This is intuitively clear when our semantics are viewed from a continuous-time perspective—any interval in which both A and B are true can be broken up into an arbitrary number of subintervals where both A and B hold. This example illustrates that there can be infinite descending chains of AMA formulas where the entire chain subsumes a given formula (but no member is equivalent to the given formula). In general, any AMA formula involving only the propositions A and B will subsume Φ .*

4.2 Motivation for AMA

MA timelines are a very natural way to capture “stretchable” sequences of state constraints. But why consider the conjunction of such sequences, i.e., AMA? We have several reasons for this language enrichment. First of all, we show below that the AMA least-general generalization (LGG) is unique; this is not true for MA. Second, and more informally, we argue that parallel conjunctive constraints can be important to learning efficiency. In particular, the space of MA formulas of length k grows in size exponentially with k , making it difficult to induce long MA formulas. However, finding several shorter MA timelines that each characterize *part* of a long sequence of changes is exponentially easier. (At least, the space to search is exponentially smaller.) The AMA conjunction of these timelines places these shorter constraints simultaneously and often captures a great deal of the concept structure. For this reason, we analyze AMA as well as MA and, in our empirical work we consider k-AMA.

The AMA language is propositional. But our intended applications are relational, or first-order, including visual event recognition. Later in this paper we show that the propositional AMA learning algorithms we develop can be effectively applied in relational domains. Our approach to first-order learning is distinctive in automatically constructing an object correspondence across examples (e.g., compare (Lavrac, Dzeroski, & Grobelnik, 1991; Roth & Yih, 2001)). Similarly, though AMA does not allow for negative state constraints we show how to obtain the practical advantages of negation, which is crucial in visual event recognition.

4.3 Conversion to First-Order Clauses

We note that AMA formulas can be translated in various ways into first-order clauses. It is not straightforward, however, to then use existing clausal generalization techniques for learning. In particular, to capture the AMA semantics in clauses, it appears necessary to define subsumption and generalization relative to a background theory that restricts us to a “continuous-time” first-order-model space. For example, consider the AMA formulas $\Phi_1 = (A \wedge B)$ and $\Phi_2 = A; B$ where A and B are propositions—from example 2 we know that $\Phi_1 \leq \Phi_2$. Now, consider a straightforward clausal translation of these formulas giving $C_1 = A(I) \wedge B(I)$ and $C_2 = A(I_1) \wedge B(I_2) \wedge \text{MEETS}(I_1, I_2) \wedge \text{SPAN}(I_1, I_2, I)$, where the I and I_j are variables that represent time intervals, *MEETS* indicates that two time intervals meet each other, and *SPAN* indicates that the union of the first two time-interval arguments equals the third time-interval argument. The intention is for satisfying assignments for I in C_1 and C_2 to indicated intervals over which Φ_1 and Φ_2 are satisfied, respectively. It should be clear that, contrary to what we want, $C_1 \not\leq C_2$ (i.e., $C_1 \not\rightarrow C_2$), since it is easy to find “unintended”

first-order models that satisfy C_1 , but not C_2 . Thus such a translation, and other similar translations, do not capture the continuous-time nature of the AMA semantics.

In order to capture the AMA semantics in a clausal setting, one might define a first-order theory that restricts us to “continuous-time” models—for example, allowing for the derivation “if property B holds over an interval, then that property also holds over all sub-intervals”. Given such a theory Σ , we have that $\Sigma \models C_1 \rightarrow C_2$, as desired. However, it is well known that least-general generalizations relative to such background theories need not exist (Plotkin, 1971), so prior work on clausal generalization does not simply subsume our results for the AMA language.

We note that for a particular training set, it may be possible to compile a “continuous-time” background theory, Σ , into a finite but “adequate” set of ground facts. Relative to such ground theories, clausal LGGs are known to always exist and thus could be used for our application. However, the only such compiling approaches that look promising to us require exploiting an analysis similar to the one given in this paper—i.e., understanding the AMA generalization and subsumption problem separately from clausal generalization, and exploiting that understanding in compiling the background theory. We have not pursued such compilations further.

Even if we are given such a compilation procedure, there are other problems with using existing clausal generalization techniques for learn AMA formulas. For the clausal translations of AMA we have found, the resulting generalizations typically fall outside of the (clausal translations of formulas in the) AMA language, so that the language bias of AMA is lost. In preliminary empirical work in our video-event recognition domain using clausal inductive logic programming (ILP) systems, we found that the learner appeared to lack the necessary language bias to find effective event definitions. While we believe it would be possible to find ways to build this language bias into ILP systems, we chose instead to define and learn within the desired language bias directly, by defining the class of AMA formulas, and studying the generalization operation on that class.

4.4 Basic Concepts and Properties of AMA

We use the following convention in naming our results: “propositions” and “theorems” are the key results of our work, with theorems being those results of the most technical difficulty, and “lemmas” are technical results needed for the later proofs of propositions or theorems. We number all the results in one sequence, regardless of type. Proofs of theorems and propositions are provided in the main text—omitted proofs of lemmas are provided in the appendix.

We give pseudo-code for our methods in a non-deterministic style. In a non-deterministic language functions can return more than one value “non-deterministically”, either because they contain non-deterministic choice points, or because they call other non-deterministic functions. Since a non-deterministic function can return more than one possible value, depending on the choices made at the choice-points encountered, specifying such a function is a natural way to specify a richly structured set (if the function has no arguments) or relation (if the function has arguments). To actually enumerate the values of the set (or the relation, once arguments are provided) one simply has to add a standard backtracking search over the different possible computations corresponding to different choices at the choice points.

4.4.1 SUBSUMPTION AND GENERALIZATION FOR STATES

The most basic formulas we deal with are states (conjunctions of propositions)—in our propositional setting computing subsumption and generalization at the state level is straightforward. A state S_1 subsumes S_2 ($S_2 \leq S_1$) iff S_1 is a subset of S_2 , viewing states as sets of propositions. From this we derive that the intersection of states is the least-general subsumer of those states and that the union of states is likewise the most general subsumee.

4.4.2 INTERDIGITATIONS

Given a set of MA timelines, we need to consider the different ways in which a model could simultaneously satisfy the timelines in the set. At the start of such a model (i.e., the first time point), the initial state from each timeline must be satisfied. At some time point in the model, one or more of the timelines can transition so that the second state in those timelines must be satisfied in place of the initial state, while the initial state of the other timelines remains satisfied. After a sequence of such transitions in subsets of the timelines, the final state of each timeline holds. Each way of choosing the transition sequence constitutes a different “interdigitation” of the timelines.

Alternatively viewed, each model simultaneously satisfying the timelines induces a *co-occurrence relation* on tuples of timeline states, one from each timeline, identifying which tuples co-occur at some point in the model. We represent this concept formally as a set of tuples of co-occurring states, i.e., a co-occurrence relation. We sometimes think of this set of tuples as ordered by the sequence of transitions. Intuitively, the tuples in an interdigitation represent the maximal time intervals over which no MA timeline has a transition, giving the co-occurring states for each such time interval.

A relation R on $X_1 \times \dots \times X_n$ is *simultaneously consistent* with orderings \leq_1, \dots, \leq_n , if, whenever $R(x_1, \dots, x_n)$ and $R(x'_1, \dots, x'_n)$, either $x_i \leq_i x'_i$, for all i , or $x'_i \leq_i x_i$, for all i . We say R is *piecewise total* if the projection of R onto each component is total—i.e., every state in any X_i appears in R .

Definition 1 *An interdigitation I of a set of MA timelines $\{\Phi_1, \dots, \Phi_n\}$ is a co-occurrence relation over $\Phi_1 \times \dots \times \Phi_n$ (viewing timelines as sets of states³) that is piecewise total, and simultaneously consistent with the state orderings of the Φ_i . We say that two states $s \in \Phi_i$ and $s' \in \Phi_j$ for $i \neq j$ co-occur in I iff some tuple of I contains both s and s' . We sometimes refer to I as a sequence of tuples, meaning the sequence lexicographically ordered by the Φ_i state orderings.*

We note that there are exponentially many interdigitations of even two MA timelines (relative to the total number of states in the timelines). Example 3 below shows an interdigitation of two MA timelines, and pseudo-code for non-deterministically generating an arbitrary interdigitation for a set of MA timelines can be found in Figure 3. Given an interdigitation I of the timelines $s_1; s_2; \dots; s_m$ and $t_1; t_2; \dots; t_n$ (and possibly others, the following basic properties of interdigitations are easily verifiable:

1. For $i < j$, if s_i and t_k co-occur in I then for all $k' < k$, s_j does not co-occur with $t_{k'}$ in I , and
2. $I(s_1, t_1)$ and $I(s_m, t_n)$.

We first use interdigitations to syntactically characterize subsumption between MA timelines. An interdigitation I of two MA timelines Φ_1 and Φ_2 is a *witness* to $\Phi_1 \leq \Phi_2$ if, for every pair of co-occurring states $s_1 \in \Phi_1$ and $s_2 \in \Phi_2$, we have $s_1 \leq s_2$. The following lemma and proposition establish the equivalence between witnessing interdigitations and MA subsumption.

Lemma 1 *For any MA timeline Φ and any model \mathcal{M} , if \mathcal{M} satisfies Φ then there is a witnessing interdigitation for $\text{MAP}(\mathcal{M}) \leq \Phi$.*

Proposition 2 *For MA timelines Φ_1 and Φ_2 , $\Phi_1 \leq \Phi_2$ iff there is an interdigitation that witnesses $\Phi_1 \leq \Phi_2$.*

Proof: We show the backward direction by induction on the number of states n in timeline Φ_1 . If $n = 1$, then the existence of a witnessing interdigitation for $\Phi_1 \leq \Phi_2$ implies that every state in Φ_2 is a subset of the single state in Φ_1 , and thus that any model of Φ_1 is a model of Φ_2 so that $\Phi_1 \leq \Phi_2$. Now, suppose for induction that the backward direction of the theorem holds whenever Φ_1 has n or fewer states. Given an arbitrary model \mathcal{M} of an $n + 1$ state Φ_1 and an interdigitation W that witnesses $\Phi_1 \leq \Phi_2$, we must show that \mathcal{M} is also a model of Φ_2 to conclude $\Phi_1 \leq \Phi_2$, as desired.

3. Recall, that, formally, MA timelines are viewed as sets of state-index pairs, rather than just sets of states. We ignore this distinction in our notation, for readability purposes, treating MA timelines as though no state is duplicated.

```

1:  an-interdigitation( $\{\Phi_1, \Phi_2, \dots, \Phi_n\}$ )
2:      // The  $\Phi_i$  are MA timelines.
3:      // Outputs an interdigitation of the  $\Phi_i$ .
4:       $S_0 = \langle \text{head}(\Phi_1), \dots, \text{head}(\Phi_n) \rangle$ ;
5:      IF for all  $1 \leq i \leq n$ ,  $|\Phi_i| = 1$ 
6:          THEN RETURN  $\langle S_0 \rangle$ ;
7:       $T' = \{\Phi_i \text{ such that } |\Phi_i| > 1\}$ ;
8:       $T'' = \text{a-non-empty-subset-of}(T')$ ;
9:      FOR i = 1 to n
10:         IF  $\Phi_i \in T''$ 
12:            THEN  $\Phi'_i = \text{rest}(\Phi_i)$ ;
12:            ELSE  $\Phi'_i = \Phi_i$ ;
13:      RETURN extend-tuple( $S_0$ , an-interdigitation( $\{\Phi'_1, \dots, \Phi'_n\}$ ));
    
```

Figure 3: Pseudo-code for `an-interdigitation()`, which non-deterministically computes an interdigitation for a set $\{\Phi_1, \dots, \Phi_n\}$ of MA timelines. The functions `head(Φ)` and `rest(Φ)` return the first state in the timeline Φ and Φ with the first state removed, respectively, and `extend-tuple(x, I)` extends a tuple I by adding a new first element x to form a longer tuple. The function `a-non-empty-subset-of(S)` non-deterministically returns an arbitrary non-empty subset of S .

Write Φ_1 as $s_1; \dots; s_{n+1}$ and Φ_2 as $t_1; \dots; t_m$. As a witnessing interdigitation, W must identify some maximal prefix $t_1; \dots; t_{m'}$ of Φ_2 made up of states that co-occur with s_1 and thus that are subsets of s_1 . Since $\mathcal{M} = \langle M, [t, t'] \rangle$ satisfies Φ_1 , by definition there must exist a $t'' \in [t, t']$ such that $\langle M, [t, t''] \rangle$ satisfies s_1 (and thus $t_1; \dots; t_{m'}$) and $\langle M, I' \rangle$ satisfies $s_2; \dots; s_{n+1}$ for I' equal to either $[t'', t']$ or $[t'' + 1, t']$. In either case, it is straightforward to construct from W a witnessing interdigitation for $s_2; \dots; s_{n+1} \leq t_{m'+1}; \dots; t_m$ and use the induction hypothesis to then show that $\langle M, I' \rangle$ must satisfy $t_{m'+1}; \dots; t_m$. It follows that \mathcal{M} satisfies Φ_2 as desired.

For the forward direction, assume that $\Phi_1 \leq \Phi_2$, and let \mathcal{M} be any model such that $\Phi_1 = \text{MAP}(\mathcal{M})$. It is clear that such \mathcal{M} exists and satisfies Φ_1 . It follows that \mathcal{M} satisfies Φ_2 , and Lemma 1 then implies that there is a witnessing interdigitation for $\text{MAP}(\mathcal{M}) \leq \Phi_2$ and thus for $\Phi_1 \leq \Phi_2$. \square

4.4.3 LEAST-GENERAL COVERING FORMULA

A logic can discriminate two models if it contains a formula that satisfies one but not the other. It turns out that AMA formulas can discriminate two models exactly when much richer *internal positive event logic (IPEL)* formulas can do so. Internal formulas are those that define event occurrence only in terms of properties within the defining interval (i.e., satisfaction by $\langle M, I \rangle$ depends only on the proposition truth values given by M inside the interval I)—positive formulas are those that do not contain negation. Appendix A gives the full syntax and semantics of IPEL (which are used only to state and prove Lemma 3). The fact that AMA can discriminate models as well as IPEL indicates that our restriction to AMA formulas retains substantial expressive power and leads to the following result which serves as the least-general covering formula (LGCF) component of our specific-to-general learning procedure. First, we introduce the concept of *model embedding*. We say that model \mathcal{M} embeds model \mathcal{M}' iff $\text{MAP}(\mathcal{M}) \leq \text{MAP}(\mathcal{M}')$.

Lemma 3 *For any $E \in \text{IPEL}$, if model \mathcal{M} embeds any model that satisfies E , then \mathcal{M} satisfies E .*

Proposition 4 *The MA-projection of a model is its LGCF for internal positive event logic (and hence for AMA), up to semantic equivalence.*

Proof: Consider model \mathcal{M} . We know that $\text{MAP}(\mathcal{M})$ covers \mathcal{M} , so it remains to show that $\text{MAP}(\mathcal{M})$ is the least general formula to do so, up to semantic equivalence.

Let E be any IPEL formula that covers \mathcal{M} . Let \mathcal{M}' be any model that is covered by $\text{MAP}(\mathcal{M})$ —we want to show that E also covers \mathcal{M}' . We know from Lemma 1 that there is a witnessing interdigitation for $\text{MAP}(\mathcal{M}') \leq \text{MAP}(\mathcal{M})$ —thus, by Proposition 2, $\text{MAP}(\mathcal{M}') \leq \text{MAP}(\mathcal{M})$ showing that \mathcal{M}' embeds \mathcal{M} . Combining these facts with Lemma 3 it follows that E also covers \mathcal{M}' and hence $\text{MAP}(\mathcal{M}) \leq E$. \square

Proposition 4 tells us that for IPEL the LGCF of a model exists, is unique, and is an MA timeline. Given this property, when an AMA formula Ψ covers all the MA timelines covered by another AMA formula Ψ' , we have $\Psi' \leq \Psi$. Thus, for the remainder of this paper when considering subsumption between formulas we can abstract away from temporal models and deal rather with MA timelines. Proposition 4 also tells us that we can compute the LGCF of a model by constructing the MA-projection of that model. Based on the definition of MA-projection, it is straightforward to derive an LGCF algorithm which runs in time polynomial in the size of the model⁴. We note that the MA-projection may contain repeated states—in practice we remove repeated states which does not change the meaning of the resulting formula (as demonstrated in Example 1).

4.4.4 COMBINING INTERDIGITATION WITH GENERALIZATION OR SPECIALIZATION

Interdigitations are useful in analyzing both conjunctions and disjunctions of MA timelines. When conjoining a set of timelines, any model of the conjunction induces an interdigitation of the timelines such that co-occurring states simultaneously hold in the model at some point (viewing states as sets, the union of the co-occurring states must hold). By constructing an interdigitation and taking the union of each tuple of co-occurring states to get a sequence of states, we get an MA timeline that forces the conjunction of the timelines to hold. We call such a sequence an “interdigitation specialization” of the timelines. Dually, an “interdigitation generalization” involving intersections of states gives an MA timeline that holds whenever the disjunction of a set of timelines holds.

Definition 2 An interdigitation generalization (specialization) of a set Σ of MA timelines is an MA timeline $s_1; \dots; s_m$, such that, for some interdigitation I of Σ with m tuples, s_j is the intersection (respectively, union) of the components of the j 'th tuple of the sequence I . The set of interdigitation generalizations (respectively, specializations) of Σ is called $\text{IG}(\Sigma)$ (respectively, $\text{IS}(\Sigma)$).

Example 3 Suppose s_1, s_2, s_3, t_1, t_2 , and t_3 are each sets of propositions (i.e., states). Consider the timelines $S = s_1; s_2; s_3$ and $T = t_1; t_2; t_3$. The relation

$$\{ \langle s_1, t_1 \rangle, \langle s_2, t_1 \rangle, \langle s_3, t_2 \rangle, \langle s_3, t_3 \rangle \}$$

is an interdigitation of S and T in which states s_1 and s_2 co-occur with t_1 , and s_3 co-occurs with t_2 and t_3 . The corresponding IG and IS members are

$$\begin{aligned} s_1 \cap t_1; s_2 \cap t_1; s_3 \cap t_2; s_3 \cap t_3 &\in \text{IG}(\{S, T\}) \\ s_1 \cup t_1; s_2 \cup t_1; s_3 \cup t_2; s_3 \cup t_3 &\in \text{IS}(\{S, T\}). \end{aligned}$$

If $t_1 \subseteq s_1, t_1 \subseteq s_2, t_2 \subseteq s_3$, and $t_3 \subseteq s_3$, then the interdigitation witnesses $S \leq T$.

Each timeline in $\text{IG}(\Sigma)$ (dually, $\text{IS}(\Sigma)$) subsumes (is subsumed by) each timeline in Σ —this is easily verified using Proposition 2. For our complexity analyses, we note that the number of states in any member of $\text{IG}(\Sigma)$ or $\text{IS}(\Sigma)$ is lower-bounded by the number of states in any of the MA timelines in Σ and is upper-bounded by the total number of states in all the MA timelines in Σ . The number of interdigitations of Σ , and thus of members of $\text{IG}(\Sigma)$ or $\text{IS}(\Sigma)$, is exponential in that same total number of states. The algorithms we present later for computing LGGs require the computation of both $\text{IG}(\Sigma)$ and $\text{IS}(\Sigma)$. Here we give pseudo-code to compute these quantities—figure 4 gives pseudo-code for the function `an-IG-member` that

4. We take the size of a model $\mathcal{M} = \langle M, I \rangle$ to be the sum over $x \in I$ of the number of true propositions in $M(x)$

```

an-IG-member( $\{\Phi_1, \Phi_2, \dots, \Phi_n\}$ )
  // The  $\Phi_i$  are MA timelines.
  // Outputs a member of  $IG(\{\Phi_1, \Phi_2, \dots, \Phi_n\})$ .
  RETURN map (intersect-tuple, an-interdigitation ( $\{\Phi_1, \dots, \Phi_n\}$ ));

```

Figure 4: Pseudo-code for an-IG-member, which non-deterministically computes a member of $IG(T)$ where T is a sequence of MA timelines. The function $\text{intersect-tuple}(I)$ takes a tuple I of sets as its argument and returns their intersection. The higher-order function $\text{map}(f, I)$ takes a function f and a tuple I as arguments and returns a tuple of the same length as I obtained by applying f to each element of I and making a tuple of the results.

non-deterministically computes an arbitrary member of $IG(\Sigma)$ (an-IS-member is identical only we replace intersection by union). Given a set of MA timelines Σ we can compute $IG(\Sigma)$ by executing all possible deterministic computation paths of the function call $\text{an-IG-member}(\Sigma)$, i.e., computing the set of results obtainable from the non-deterministic function for all possible decisions at non-deterministic choice-points.

We now give a useful lemma and a proposition concerning the relationships between conjunctions and disjunctions of MA concepts (the former being AMA concepts). For convenience here, we use disjunction on MA concepts, producing formulas outside of AMA with the obvious interpretation.

Lemma 5 *Given an MA formula Φ that subsumes each member of a set Σ of MA formulas, Φ also subsumes some member Φ' of $IG(\Sigma)$. Dually, when Φ is subsumed by each member of Σ , we have that Φ is also subsumed by some member Φ' of $IS(\Sigma)$. In each case, the length of Φ' is bounded by the size of Σ .*

Proposition 6 *The following hold:*

1. (and-to-or) *The conjunction of a set Σ of MA timelines equals the disjunction of the timelines in $IS(\Sigma)$.*
2. (or-to-and) *The disjunction of a set Σ of MA timelines is subsumed by the conjunction of the timelines in $IG(\Sigma)$.*

Proof: To prove part 2 recall that for any $\Phi \in \Sigma$ and any $\Phi' \in IG(\Sigma)$ we have that $\Phi \leq \Phi'$. From this it is immediate that $(\bigvee \Sigma) \leq (\bigwedge IG(\Sigma))$. Using a dual argument we can show that $(\bigvee IS(\Sigma)) \leq (\bigwedge \Sigma)$. It remains to show that $(\bigwedge \Sigma) \leq (\bigvee IS(\Sigma))$, which is equivalent to showing that any timeline subsumed by $(\bigwedge \Sigma)$ is also subsumed by $(\bigvee IS(\Sigma))$ (by Proposition 4). Consider any MA timeline Φ such that $\Phi \leq (\bigwedge \Sigma)$ —this implies that each member of Σ subsumes Φ . Lemma 5 then implies that there is some $\Phi' \in IS(\Sigma)$ such that $\Phi \leq \Phi'$. From this we get that $\Phi \leq (\bigvee IS(\Sigma))$, as desired. \square

Using “and-to-or”, we can now reduce AMA subsumption to MA subsumption, with an exponential increase in the problem size.

Proposition 7 *For AMA Ψ_1 and Ψ_2 , $(\Psi_1 \leq \Psi_2)$ if and only if for all $\Phi_1 \in IS(\Psi_1)$ and $\Phi_2 \in \Psi_2$, $\Phi_1 \leq \Phi_2$*

Proof: For the forward direction we show the contrapositive. Assume there is a $\Phi_1 \in IS(\Psi_1)$ and a $\Phi_2 \in \Psi_2$ such that $\Phi_1 \not\leq \Phi_2$. Thus, there is an MA timeline Φ such that $\Phi \leq \Phi_1$ but $\Phi \not\leq \Phi_2$. This tells us that $\Phi \leq (\bigvee IS(\Psi_1))$ and that $\Phi \not\leq \Psi_2$, thus $(\bigvee IS(\Psi_1)) \not\leq \Psi_2$ and by “and-to-or” we get that $\Psi_1 \not\leq \Psi_2$.

For the backward direction assume that for all $\Phi_1 \in IS(\Psi_1)$ and $\Phi_2 \in \Psi_2$ that $\Phi_1 \leq \Phi_2$. This tells us that for each $\Phi_1 \in IS(\Psi_1)$ that $\Phi_1 \leq \Psi_2$ —thus, $\Psi_1 = (\bigvee IS(\Psi_1)) \leq \Psi_2$. \square

```

MA-subsumes ( $\Phi_1, \Phi_2$ )
// Input:  $\Phi_1 = s_1; \dots; s_m$  and  $\Phi_2 = t_1; \dots; t_n$ 
// Output:  $\Phi_1 \leq \Phi_2$ 

1. IF there is a path from  $v_{1,1}$  to  $v_{m,n}$  in  $SG(\Phi_1, \Phi_2)$  THEN RETURN TRUE. For example,
    (a) Create an array  $\text{Reachable}(i,j)$  of boolean values, all FALSE, for  $0 \leq i \leq m$  and  $0 \leq j \leq n$ .
    (b) FOR  $i = 1$  to  $m$ ,  $\text{Reachable}(i, 0) = \text{TRUE}$ ;
        FOR  $j = 1$  to  $n$ ,  $\text{Reachable}(0, j) = \text{TRUE}$ ;
        FOR  $i = 1$  to  $m$ 
            FOR  $j = 1$  to  $n$ 
                 $\text{Reachable}(i, j) = (t_i \subseteq s_j \wedge (\text{Reachable}(i-1, j) \vee$ 
                     $\text{Reachable}(i, j-1) \vee$ 
                     $\text{Reachable}(i-1, j-1))$ );
    (c) IF  $\text{Reachable}(m, n)$  THEN RETURN TRUE;

2. Otherwise, RETURN FALSE;
    
```

Figure 5: Pseudo-code for the MA subsumption algorithm. $SG(\Phi_1, \Phi_2)$ is the subsumption graph defined in the main text.

5. Subsumption and Generalization

In this section we study subsumption and generalization of AMA formulas. First, we give a polynomial-time algorithm for deciding subsumption between MA formulas and then show that deciding subsumption for AMA formulas is coNP-complete. Second we give algorithms and complexity bounds for the construction of least-general generalization (LGG) formulas based on our analysis of subsumption, including existence, uniqueness, lower/upper bounds, and an algorithm for the LGG on AMA formulas. Third, we introduce a polynomial-time-computable syntactic notion of subsumption and an algorithm that computes the corresponding syntactic LGG that is exponentially faster than our semantic LGG algorithm. Fourth, in Section 5.4, we give a detailed example showing the steps performed by our LGG algorithms to compute the semantic and syntactic LGGs of two AMA formulas.

5.1 Subsumption

All our methods rely critically on a novel algorithm for deciding the subsumption question $\Phi_1 \leq \Phi_2$ between MA formulas Φ_1 and Φ_2 in polynomial-time. We note that merely searching the possible interdigitations of Φ_1 and Φ_2 for a witnessing interdigitation provides an obvious decision procedure for the subsumption question—however, there are, in general, exponentially many such interdigitations. We reduce the MA subsumption problem to finding a path in a graph on pairs of states in $\Phi_1 \times \Phi_2$, a polynomial-time operation. Pseudo-code for the resulting MA subsumption algorithm is shown in Figure 5. The main data structure used by the MA subsumption algorithm is the subsumption graph.

Definition 3 *The subsumption graph of two MA timelines $\Phi_1 = s_1; \dots; s_m$ and $\Phi_2 = t_1; \dots; t_n$ (written $SG(\Phi_1, \Phi_2)$) is a directed graph $G = \langle V, E \rangle$ with vertex set V equal to $\{v_{i,j} \mid 1 \leq i \leq m, 1 \leq j \leq n\}$. The (directed) edge set E equals $\{\langle v_{i,j}, v_{i',j'} \rangle \mid s_i \leq t_j, s_{i'} \leq t_{j'}, i \leq i' \leq i+1, j \leq j' \leq j+1\}$.*

To achieve a polynomial-time bound one can simply use any polynomial-time pathfinding algorithm. In our case the special structure of the subsumption graph can be exploited to determine if the desired path exists in $O(mn)$ time, as the example method shown in the pseudo-code illustrates. The following theorem asserts the correctness of the algorithm assuming a correct polynomial-time path-finding method is used.

Lemma 8 *Given MA timelines $\Phi_1 = s_1; \dots; s_m$ and $\Phi_2 = t_1; \dots; t_n$, there is a witnessing interdigitation for $\Phi_1 \leq \Phi_2$ iff there is a path in the subsumption graph $SG(\Phi_1, \Phi_2)$ from $v_{1,1}$ to $v_{m,n}$.*

Theorem 9 *Given MA timelines Φ_1 and Φ_2 , $MA\text{-subsumes}(\Phi_1, \Phi_2)$ decides $\Phi_1 \leq \Phi_2$ in polynomial time.*

Proof: The algorithm clearly runs in polynomial time. Lemma 8 tells us that line 2 of the algorithm will return TRUE iff there is a witnessing interdigitation. Combining this with Proposition 2 shows that the algorithm returns TRUE iff $\Phi_1 \leq \Phi_2$. \square

Given this polynomial-time algorithm for MA subsumption, Proposition 7 immediately suggests an exponential-time algorithm for deciding AMA subsumption—by computing MA subsumption between the exponentially many IS timelines of one formula and the timelines of the other formula. Our next theorem suggests that we cannot do any better than this in the worst case—we argue that AMA subsumption is coNP-complete by reduction from boolean satisfiability. Readers uninterested in the technical details of this argument may skip directly to Section 5.2.

To develop a correspondence between boolean satisfiability problems and AMA formulas, which lack negation, we imagine that each boolean variable has two AMA propositions, one for “true” and one for “false”. In particular, given a boolean satisfiability problem over n variables p_1, \dots, p_n , we take the set $PROP_n$ to be the set containing $2n$ AMA propositions $True_k$ and $False_k$ for each k between 1 and n . We can now represent a truth assignment A to the p_i variables with an AMA state s_A given as follows:

$$s_A = \{True_i \mid 1 \leq i \leq n, A(p_i) = \text{true}\} \cup \{False_i \mid 1 \leq i \leq n, A(p_i) = \text{false}\}$$

As Proposition 7 suggests, checking AMA subsumption critically involves the exponentially many interdigitation specializations of the timelines of one of the AMA formulas. In our proof, we design an AMA formula whose interdigitation specializations can be seen to correspond to truth assignments⁵ to boolean variables, as shown in the following lemma.

Lemma 10 *Given some n , let Ψ be the conjunction of the timelines*

$$\bigcup_{i=1}^n \{ (PROP_n; True_i; False_i; PROP_n), (PROP_n; False_i; True_i; PROP_n) \}.$$

We have the following facts about truth assignments to the Boolean variables p_1, \dots, p_n :

1. *For any truth assignment A , $PROP_n; s_A; PROP_n$ is semantically equivalent to some member of $IS(\Psi)$.*
2. *For each $\Phi \in IS(\Psi)$ there is a truth assignment A such that $\Phi \leq PROP_n; s_A; PROP_n$.*

With this lemma in hand, we can now tackle the complexity of AMA subsumption.

Theorem 11 *Deciding AMA subsumption is coNP-complete.*

Proof: We first show that deciding the AMA-subsumption of Ψ_1 by Ψ_2 is in coNP by providing a polynomial-length certificate for any “no” answer. This certificate for non-subsumption is an interdigitation of the timelines of Ψ_1 that yields a member of $IS(\Psi_1)$ not subsumed by Ψ_2 . Such a certificate can be checked in polynomial time: given such an interdigitation, the corresponding member of $IS(\Psi_1)$ can be computed in time polynomial in the size of Ψ_1 , and we can then test whether the resulting timeline is subsumed by each timeline in Ψ_2 using the polynomial-time MA-subsumption algorithm. Proposition 7 guarantees that $\Psi_1 \not\leq \Psi_2$ iff there is a timeline in $IS(\Psi_1)$ that is not subsumed by every timeline in Ψ_2 , so that such a certificate will exist exactly when the answer to a subsumption query is “no”.

To show coNP-hardness we reduce the problem of deciding the satisfiability of a 3-SAT formula $S = C_1 \wedge \dots \wedge C_m$ to the problem of recognizing non-subsumption between AMA formulas. Here, each C_i is

5. A truth assignment is a function mapping boolean variables to true or false.

$(l_{i,1} \vee l_{i,2} \vee l_{i,3})$ and each $l_{i,j}$ either a proposition p chosen from $P = \{p_1, \dots, p_n\}$ or its negation $\neg p$. The idea of the reduction is to construct an AMA formula Ψ for which we view the exponentially many members of $\text{IS}(\Psi)$ as representing truth assignments. We then construct an MA timeline Φ that we view as representing $\neg S$ and show that S is satisfiable iff $\Psi \not\leq \Phi$.

Let Ψ be as defined in Lemma 10. Let Φ be the formula $s_1; \dots; s_m$, where

$$s_i = \{\text{False}_j \mid l_{i,k} = p_j \text{ for some } k\} \cup \{\text{True}_j \mid l_{i,k} = \neg p_j \text{ for some } k\}.$$

Each s_i can be thought of as asserting “not C_i ”. We start by showing that if S is satisfiable then $\Psi \not\leq \Phi$. Assume that S is satisfied via a truth assignment A —we know from Lemma 10 that there is a $\Phi' \in \text{IS}(\Psi)$ that is semantically equivalent to $\text{PROP}_n; s_A; \text{PROP}_n$. We show that $\text{PROP}_n; s_A; \text{PROP}_n$ is not subsumed by Φ , to conclude $\Psi \not\leq \Phi$ using Proposition 7, as desired. Suppose for contradiction that $\text{PROP}_n; s_A; \text{PROP}_n$ is subsumed by Φ —then the state s_A must be subsumed by some state s_i in Φ . Consider the corresponding clause C_i of S . Since A satisfies S we have that C_i is satisfied and at least one of its literals $l_{i,k}$ must be true. Assume that $l_{i,k} = p_j$ (a dual argument holds for $l_{i,k} = \neg p_j$), then we have that s_i contains False_j while s_A contains True_j but not False_j —thus, we have that $s_A \not\leq s_i$ (since $s_i \not\leq s_A$), contradicting our choice of i .

To complete the proof, we now assume that S is unsatisfiable and show that $\Psi \leq \Phi$. Using Proposition 7, we consider arbitrary Φ' in $\text{IS}(\Psi)$ —we will show that $\Phi' \leq \Phi$. From Lemma 10 we know there is some truth assignment A such that $\Phi' \leq \text{PROP}_n; s_A; \text{PROP}_n$. Since S is unsatisfiable we know that some C_i is not satisfied by A and hence $\neg C_i$ is satisfied by A . This implies that each primitive proposition in s_i is in s_A . Let W be the following interdigitation between $T = \text{PROP}_n; s_A; \text{PROP}_n$ and $\Phi = s_1; \dots; s_m$:

$$W = \{ \langle \text{PROP}_n, s_1 \rangle \langle \text{PROP}_n, s_2 \rangle \cdots \langle \text{PROP}_n, s_i \rangle \langle s_A, s_i \rangle \langle \text{PROP}_n, s_i \rangle \langle \text{PROP}_n, s_{i+1} \rangle \cdots \langle \text{PROP}_n, s_m \rangle \}$$

We see that in each tuple of co-occurring states given above that the state from T is subsumed by the state from Φ . Thus W is a witnessing interdigitation to $\text{PROP}_n; s_A; \text{PROP}_n \leq \Phi$, which then holds by Proposition 2—combining this with $\Phi' \leq \text{PROP}_n; s_A; \text{PROP}_n$ we get that $\Phi' \leq \Phi$, as desired. \square

Given this hardness result we later define a weaker polynomial-time-computable subsumption notion for use in our learning algorithms.

5.2 Least-General Generalization.

The existence of an AMA LGG is nontrivial as there can be infinite chains of increasingly specific formulas all of which generalize given formulas. Example 2 demonstrated such chains for an MA subsumee and can be extended for AMA subsumees. For example, each member of the chain $P; Q, P; Q; P; Q, P; Q; P; Q; P; Q, \dots$ covers $\Psi_1 = (P \wedge Q); Q$ and $\Psi_2 = P; (P \wedge Q)$. Despite such complications, the AMA LGG does exist.

Theorem 12 *There is an LGG for any finite set Σ of AMA formulas that is subsumed by all other generalizations of Σ .*

Proof: Let Γ be the set $\bigcup_{\Psi' \in \Sigma} \text{IS}(\Psi')$. Let Ψ be the conjunction of all the MA timelines that generalize Γ while having size no larger than Γ . Since there are only a finite number of primitive propositions, there are only a finite number of such timelines, so Ψ is well defined⁶. We show that Ψ is a least-general generalization of Σ . First, note that each timeline in Ψ generalizes Γ and thus Σ (by Proposition 6), so Ψ must generalize Σ . Now, consider arbitrary generalization Ψ' of Σ . Proposition 7 implies that Ψ' must generalize each formula in Γ . Lemma 5 then implies that each timeline of Ψ' must subsume a timeline Φ that is no longer than the size of Γ and that also subsumes the timelines of Γ . But then Φ must be a timeline of Ψ , by our choice of Ψ ,

6. There must be at least one such timeline, the timeline with one state containing all propositions

```

1: semantic-LGG( $\{\Psi_1, \Psi_2, \dots, \Psi_m\}$ )
2:   // The  $\Psi_i$  are AMA formulas.
3:   // Outputs an LGG of the  $\Psi_i$ .
4:    $S = \{\}$ ;
5:   For  $i = 1$  to  $m$ 
6:     For each  $\Phi$  in all-values(an-IS-member( $\Psi_i$ ))
7:       IF ( $\forall \Phi' \in S . \Phi \not\leq \Phi'$ )
8:         THEN  $S' = \{\Phi'' \in S \mid \Phi'' \leq \Phi\}$ ;
9:          $S = (S - S') \cup \{\Phi\}$ ;
10:   $G = \{\}$ ;
11:  For each  $\Phi$  in all-values(an-IG-member( $S$ ))
12:    IF ( $\forall \Phi' \in G . \Phi' \not\leq \Phi$ )
13:      THEN  $G' = \{\Phi'' \in G \mid \Phi \leq \Phi''\}$ ;
14:       $G = (G - G') \cup \{\Phi\}$ ;
15:  RETURN ( $\bigwedge G$ )
    
```

Figure 6: Pseudo-code for computing the semantic AMA LGG of a set of AMA formulas.

so that every timeline of Ψ' subsumes a timeline of Ψ . It follows that Ψ' subsumes Ψ , and that Ψ is an LGG of Σ subsumed by all other LGGs of Σ , as desired. \square

Given that the AMA LGG exists and is unique we now show how to compute it. Our first step is to strengthen “or-to-and” from Proposition 6 to get an LGG for the MA sublanguage.

Theorem 13 *For a set Σ of MA formulas, the conjunction of all MA timelines in $\text{IG}(\Sigma)$ is an AMA LGG of Σ .*

Proof: Let Ψ be the specified conjunction. Since each timeline of $\text{IG}(\Sigma)$ subsumes all timelines in Σ , Ψ subsumes each member of Σ . To show Ψ is a least-general such formula, consider an AMA formula Ψ' that also subsumes all members of Σ . Since each timeline of Ψ' must subsume all members of Σ , Lemma 5 implies that each timeline of Ψ' subsumes a member of $\text{IG}(\Sigma)$ and thus each timeline of Ψ' subsumes Ψ . This implies $\Psi \leq \Psi'$. \square

We can now characterize the AMA LGG using IS and IG.

Theorem 14 *$\text{IG}(\bigcup_{\Psi \in \Sigma} \text{IS}(\Psi))$ is an AMA LGG of the set Σ of AMA formulas.*

Proof: Let $\Sigma = \{\Psi_1, \dots, \Psi_n\}$ and $E = \Psi_1 \vee \dots \vee \Psi_n$. We know that the AMA LGG of Σ must subsume E , or it would fail to subsume one of the Ψ_i . Using “and-to-or” we can represent E as a disjunction of MA timelines given by $E = (\bigvee \text{IS}(\Psi_1)) \vee \dots \vee (\bigvee \text{IS}(\Psi_n))$. Any AMA LGG must be a least-general formula that subsumes E —i.e., an AMA LGG of the set of MA timelines $\bigcup\{\text{IS}(\Psi) \mid \Psi \in \Sigma\}$. Theorem 13 tells us that an LGG of these timelines is given by $\text{IG}(\bigcup\{\text{IS}(\Psi) \mid \Psi \in \Sigma\})$. \square

Theorem 14 leads directly to an algorithm for computing the AMA LGG—figure 6 gives pseudo-code for the computation. Lines 4-9 of the pseudo-code correspond to the computation of $\bigcup\{\text{IS}(\Psi) \mid \Psi \in \Sigma\}$, where timelines are not included in the set if they are subsumed by timelines already in the set (which can be checked with the polynomial time MA subsumption algorithm). This pruning, accomplished by the IF test in line 7, often drastically reduces the size of the timeline set for which we perform the subsequent IG computation—the final result is not affected by the pruning since the subsequent IG computation is a generalization step. The remainder of the pseudo-code corresponds to the computation of $\text{IG}(\bigcup\{\text{IS}(\Psi) \mid \Psi \in \Sigma\})$ where we do

not include timelines in the final result that subsume some other timeline in the set. This pruning step (the IF test in line 12) is sound since when one timeline subsumes another, the conjunction of those timelines is equivalent to the most specific one. Section 5.4.1 traces the computations of this algorithm for an example LGG calculation.

Since the sizes of both $\text{IS}(\cdot)$ and $\text{IG}(\cdot)$ are exponential in the sizes of their inputs, the code in Figure 6 is doubly exponential in the input size. We conjecture that we cannot do better than this, but we have not yet proven a doubly exponential lower bound for the AMA case. When the input formulas are MA timelines the algorithm takes singly exponential time, since $\text{IS}(\{\Phi\}) = \Phi$ when Φ is in MA. We now prove an exponential lower bound when the input formulas are in MA. Again, readers uninterested in the technical details of this proof can safely skip forward to Section 5.3.

For this argument, we take the available primitive propositions to be those in the set $\{p_{i,j} \mid 1 \leq i \leq n, 1 \leq j \leq n\}$, and consider the MA timelines

$$\begin{aligned} \Phi_1 &= s_{1,*}; s_{2,*}; \dots; s_{n,*} \\ \text{and } \Phi_2 &= s_{*,1}; s_{*,2}; \dots; s_{*,n}, \text{ where} \\ s_{i,*} &= p_{i,1} \wedge \dots \wedge p_{i,n} \\ \text{and } s_{*,j} &= p_{1,j} \wedge \dots \wedge p_{n,j}. \end{aligned}$$

We will show that any AMA LGG of Φ_1 and Φ_2 must contain an exponential number of timelines. In particular, we will show that any AMA LGG is equivalent to the conjunction of a subset of $\text{IG}(\{\Phi_1, \Phi_2\})$, and that certain timelines may not be omitted from such a subset.

Lemma 15 *Any AMA LGG Ψ of a set Σ of MA timelines is equivalent to a conjunction Ψ' of timelines from $\text{IG}(\Sigma)$ with $|\Psi'| \leq |\Psi|$*

Proof: Lemma 5 implies that any timeline Φ in Ψ must subsume some timeline $\Phi' \in \text{IG}(\Sigma)$. But then the conjunction Ψ' of such Φ' must be equivalent to Ψ , since it clearly covers Σ and is covered by the LGG Ψ . Since Ψ' was formed by taking one timeline from $\text{IG}(\Sigma)$ for each timeline in Ψ , we have $|\Psi'| \leq |\Psi|$. \square

We can complete our argument then by showing that exponentially many timelines in $\text{IG}(\{\Phi_1, \Phi_2\})$ cannot be omitted from such a conjunction while it remains an LGG.

Notice that for any i, j we have that $s_{i,*} \cap s_{*,j} = p_{i,j}$. This implies that any state in $\text{IG}(\{\Phi_1, \Phi_2\})$ contains exactly one proposition, since each such state is formed by intersecting a state from Φ_1 and Φ_2 . Furthermore, the definition of interdigitation, applied here, implies the following two facts for any timeline $q_1; q_2; \dots; q_m$ in $\text{IG}(\{\Phi_1, \Phi_2\})$:

1. $q_1 = p_{1,1}$ and $q_m = p_{n,n}$.
2. For consecutive states $q_k = p_{i,j}$ and $q_{k+1} = p_{i',j'}$, i' is either i or $i+1$, j' is either j or $j+1$, and not both $i = i'$ and $j = j'$.

Together these facts imply that any timeline in $\text{IG}(\{\Phi_1, \Phi_2\})$ is a sequence of propositions starting with $p_{1,1}$ and ending with $p_{n,n}$ such that any consecutive propositions $p_{i,j}; p_{i',j'}$ are different with i' equal to i or $i+1$ and j' equal to j or $j+1$. We call a timeline in $\text{IG}(\{\Phi_1, \Phi_2\})$ *square* if and only if for each pair of consecutive propositions $p_{i,j}$ and $p_{i',j'}$ have either $i' = i$ or $j' = j$. The following lemma implies that no square timeline can be omitted from the conjunction of timelines in $\text{IG}(\Phi_1, \Phi_2)$ if it is to remain an LGG of Φ_1 and Φ_2 .

Lemma 16 *Let Φ_1 and Φ_2 be as given above and let $\Psi = \bigwedge \text{IG}(\{\Phi_1, \Phi_2\})$. For any Ψ' whose timelines are a subset of those in Ψ that omits some square timeline, we have $\Psi < \Psi'$.*

The number of square timelines in $\text{IG}(\{\Phi_1, \Phi_2\})$ is equal to $\frac{(2n-2)!}{(n-1)!(n-1)!}$ and hence is exponential in the size of Φ_1 and Φ_2 . We have now completed the proof of the following result.

Theorem 17 *The smallest LGG of two MA formulas can be exponentially large.*

Proof: By Lemma 15, any AMA LGG Ψ' of Φ_1 and Φ_2 is equivalent to a conjunction of the same number of timelines chosen from $\text{IG}(\{\Phi_1, \Phi_2\})$. However, by Lemma 16, any such conjunction must have at least $\frac{(2n-2)!}{(n-1)!(n-1)!}$ timelines, and then so must Ψ' , which must then be exponentially large. \square

Conjecture 18 *The smallest LGG of two AMA formulas can be doubly-exponentially large.*

We now show that our lower-bound on AMA LGG complexity is not merely a consequence of the existence of large AMA LGGs. Even when there is a small LGG, it can be expensive to compute due to the difficulty of testing AMA subsumption:

Theorem 19 *Determining whether a formula Ψ is an AMA LGG for two given AMA formulas Ψ_1 and Ψ_2 is co-NP-hard, and is in co-NEXP, in the size of all three formulas together.*

Proof: To show co-BP-hardness we use a straightforward reduction from AMA subsumption. Given two AMA formulas Ψ_1 and Ψ_2 we decide $\Psi_1 \leq \Psi_2$ by asking whether Ψ_2 is an AMA LGG of Ψ_1 and Ψ_2 . Clearly $\Psi_1 \leq \Psi_2$ iff Ψ_2 is an LGG of the two formulas.

To show the co-NEAP upper bound, note that we can check in exponential time whether $\Psi_1 \leq \Psi$ and $\Psi_2 \leq \Psi$ using Proposition 7 and the polynomial-time MA subsumption algorithm. It remains to show that we can check whether Ψ is *not* the “least” subsumer. Since Theorem 14 shows that the LGG of Ψ_1 and Ψ_2 is $\text{IG}(\text{IS}(\Psi_1) \cup \text{IS}(\Psi_2))$, if Ψ is not the LGG then $\Psi \not\leq \text{IG}(\text{IS}(\Psi_1) \cup \text{IS}(\Psi_2))$. Thus, by Proposition 7, if Ψ is not a least subsumer, there must be timelines $\Phi_1 \in \text{IS}(\Psi)$ and $\Phi_2 \in \text{IG}(\text{IS}(\Psi_1) \cup \text{IS}(\Psi_2))$ such that $\Phi_1 \not\leq \Phi_2$. We can then use exponentially long certificates for “No” answers: each certificate is a pair of an interdigitation I_1 of Ψ and an interdigitation I_2 of $\text{IS}(\Psi_1) \cup \text{IS}(\Psi_2)$, such that the corresponding members $\Phi_1 \in \text{IS}(\Psi)$ and $\Phi_2 \in \text{IG}(\text{IS}(\Psi_1) \cup \text{IS}(\Psi_2))$ have $\Phi_1 \not\leq \Phi_2$. Given the pair of certificates I_1 and I_2 , Φ_1 can be computed in polynomial time, Φ_2 can be computed in exponential time, and the subsumption between them can be checked in polynomial time (relative to their size, which can be exponential). If Ψ is the LGG then $\Psi \leq \text{IG}(\text{IS}(\Psi_1) \cup \text{IS}(\Psi_2))$, so that no such certificates will exist. \square

5.3 Syntactic Subsumption and Syntactic Least-General Generalization.

Given the intractability results for semantic AMA subsumption, we now introduce a tractable generality notion, syntactic subsumption, and discuss the corresponding LGG problem. The use of syntactic forms of generality for efficiency is familiar in ILP (Muggleton & De Raedt, 1994)—where, for example, θ -subsumption is often used in place of the entailment generality relation. Unlike AMA semantic subsumption, syntactic subsumption requires checking only polynomially many MA subsumptions, each in polynomial time (via Theorem 9).

Definition 4 *AMA Ψ_1 is syntactically subsumed by AMA Ψ_2 (written $\Psi_1 \leq_{\text{syn}} \Psi_2$) iff for each MA timeline $\Phi_2 \in \Psi_2$, there is an MA timeline $\Phi_1 \in \Psi_1$ such that $\Phi_1 \leq \Phi_2$.*

Proposition 20 *AMA syntactic subsumption can be decided in polynomial time.*

Syntactic subsumption trivially implies semantic subsumption—however, the converse does not hold in general. Consider the AMA formulas $(A; B) \wedge (B; A)$, and $A; B; A$ where A and B are primitive propositions. We have $(A; B) \wedge (B; A) \leq A; B; A$; however, we have neither $A; B \leq A; B; A$ nor $B; A \leq A; B; A$, so that $A; B; A$ does not syntactically subsume $(A; B) \wedge (B; A)$. Syntactic subsumption fails to recognize constraints that are only derived from the interaction of timelines within a formula.

Syntactic Least-General Generalization. The *syntactic AMA LGG* is the syntactically least-general AMA formula that syntactically subsumes the input AMA formulas. Here, “least” means that no formula properly syntactically subsumed by a syntactic LGG can syntactically subsume the input formulas. Based on the hardness gap between syntactic and semantic AMA subsumption, one might conjecture that a similar gap

exists between the syntactic and semantic LGG problems. Proving such a gap exists requires closing the gap between the lower and upper bounds on AMA LGG shown in Theorem 14 in favor of the upper bound, as suggested by Conjecture 18. While we cannot yet show a hardness gap between semantic and syntactic LGG, we do give a syntactic LGG algorithm that is exponentially more efficient than the best semantic LGG algorithm we have found (that of Theorem 14). First, we show that syntactic LGG’s exist and are unique up to mutual syntactic subsumption (and hence up to semantic equivalence).

Theorem 21 *There exists a syntactic LGG for any AMA formula set Σ that is syntactically subsumed by all syntactic generalizations of Σ .*

Proof: Let Ψ be the conjunction of all the MA timelines that syntactically generalize Σ while having size no larger than Σ . As in the proof of Theorem 12, Ψ is well defined. We show that Ψ is a syntactic LGG for Σ . First, note that Ψ syntactically generalizes Σ because each timeline of Ψ generalizes a timeline in every member of Σ , by the choice of Ψ . Now consider an arbitrary syntactic generalization Ψ' of Σ . By the definition of syntactic subsumption, each timeline Φ in Ψ' must subsume some timeline Φ_α in each member α of Σ . Lemma 5 then implies that there is a timeline Φ' of size no larger than Σ that subsumes all the Φ_α while being subsumed by Φ . By our choice of Ψ , the timeline Φ' must be a timeline of Ψ . It follows then that Ψ' syntactically subsumes Ψ , and that Ψ is a syntactic LGG of Σ subsumed by all other generalizations of Σ . \square

In general, we know that semantic and syntactic LGGs are different, though clearly the syntactic LGG is a semantic generalization and so must subsume the semantic LGG. For example, $(A; B) \wedge (B; A)$, and $A; B; A$ have a semantic LGG of $A; B; A$, as discussed above; but their syntactic LGG is $(A; B; \mathbf{true}) \wedge (\mathbf{true}; B; A)$, which subsumes $A; B; A$ but is not subsumed by $A; B; A$. Even so, for MA formulas:

Proposition 22 *For MA Φ and AMA Ψ , $\Phi \leq_{\text{syn}} \Psi$ is equivalent to $\Phi \leq \Psi$.*

Proof: The forward direction is immediate since we already know syntactic subsumption implies semantic subsumption. For the reverse direction, note that $\Phi \leq \Psi$ implies that each timeline of Ψ subsumes Φ —thus since Φ is a single timeline each timeline in Ψ subsumes “some timeline” in Φ which is the definition of syntactic subsumption. \square

Proposition 23 *Any syntactic AMA LGG for an MA formula set Σ is also a semantic LGG for Σ .*

Proof: Now, consider a syntactic LGG Ψ for Σ . Proposition 22 implies that Ψ is a semantic generalization of Σ . Consider any semantic LGG Ψ' of Σ . We show that $\Psi \leq \Psi'$ to conclude that Ψ is a semantic LGG for Σ . Proposition 22 implies that Ψ' syntactically subsumes Σ . It follows that $\Psi' \wedge \Psi$ syntactically subsumes Σ . But, $\Psi' \wedge \Psi$ is syntactically subsumed by Ψ , which is a syntactic LGG of Σ —it follows that $\Psi' \wedge \Psi$ syntactically subsumes Ψ , or Ψ would not be a *least* syntactic generalization of Σ . But then $\Psi \leq (\Psi' \wedge \Psi)$, which implies $\Psi \leq \Psi'$, as desired. \square

We note that the stronger result stating that a formula Ψ is a syntactic LGG of a set Σ of MA formulas if and only if it is a semantic LGG of Σ is *not* an immediate consequence of our results above. At first examination, the strengthening appears trivial, given the equivalence of $\Phi \leq \Psi$ and $\Phi \leq_{\text{syn}} \Psi$ for MA Φ . However, being semantically least is *not* necessarily a stronger condition than being syntactically least—we have not ruled out the possibility that a semantically least generalization Ψ may syntactically subsume another generalization that is semantically (but not syntactically) equivalent. (This question is open, as we have not found an example of this phenomenon either.)

Proposition 23 together with Theorem 21 have the nice consequence for our learning approach that the syntactic LGG of two AMA formulas is a semantic LGG of those formulas, as long as the original formulas are themselves syntactic LGGs of sets of MA timelines. Because our learning approach starts with training examples that are converted to MA timelines using the LGCF operation, the syntactic LGGs computed (whether combining all the training examples at once, or incrementally computing syntactic LGGs of parts of

the training data) are always syntactic LGGs of sets of MA timelines and hence are also semantic LGGs, in spite of the fact that syntactic subsumption is weaker than semantic subsumption. We note, however, that the resulting semantic LGGs may be considerably larger than the smallest corresponding semantic LGG (which may not be a syntactic LGG at all).

Using Proposition 23, we now show that we cannot hope for a polynomial-time syntactic LGG algorithm.

Theorem 24 *The smallest syntactic LGG of two MA formulas can be exponentially large.*

Proof: Suppose there is always a syntactic LGG of two MA formulas that is not exponentially large. Since by Proposition 23 each such formula is also a semantic LGG, there is always a semantic LGG of two MA formulas that is not exponentially large. This contradicts Theorem 17. \square

While this is discouraging, we have an algorithm for the syntactic LGG whose time complexity matches this lower-bound, unlike the semantic LGG case, where the best algorithm we have is doubly exponential in the worst case. Theorem 14 yields an exponential time method for computing the semantic LGG of a set of MA timelines Σ —since for a timeline Φ , $IS(\Phi) = \Phi$, we can simply conjoin all the timelines of $IG(\Sigma)$. Given a set of AMA formulas, the syntactic LGG algorithm uses this method to compute the polynomially-many semantic LGGs of sets of timelines, one chosen from each input formula, and conjoins all the results.

Theorem 25 *The formula $\bigwedge_{\Phi_i \in \Psi_i} IG(\{\Phi_1, \dots, \Phi_n\})$ is a syntactic LGG of the AMA formulas Ψ_1, \dots, Ψ_n .*

Proof: Let Ψ be $\bigwedge_{\Phi_i \in \Psi_i} IG(\{\Phi_1, \dots, \Phi_n\})$. Each timeline Φ of Ψ must subsume each Ψ_i because Φ is an output of IG on a set containing a timeline of Ψ_i —thus Ψ syntactically subsumes each Ψ_i . To show that Ψ is a syntactically least such formula, consider a Ψ' that syntactically subsumes every Ψ_i . We show that $\Psi \leq_{\text{syn}} \Psi'$ to conclude. Each timeline Φ' in Ψ' subsumes a timeline $T_i \in \Psi_i$, for each i , by our assumption that $\Psi_i \leq_{\text{syn}} \Psi'$. But then by Lemma 5, Φ' must subsume a member of $IG(\{T_1, \dots, T_n\})$ —and that member is a timeline of Ψ —so each timeline Φ' of Ψ' subsumes a timeline of Ψ . We conclude $\Psi \leq_{\text{syn}} \Psi'$, as desired. \square

This theorem yields an algorithm that computes a syntactic AMA LGG in exponential time—pseudo-code for this method is given in Figure 7. The exponential time bound follows from the fact that there are exponentially many ways to choose Φ_1, \dots, Φ_m in line 5, and for each of these there are exponentially many semantic-LGG members in line 6 (since the Φ_i are all MA timelines)—the product of these two exponentials is still an exponential.

The formula returned by the algorithm shown is actually a subset of the syntactic LGG given by Theorem 25. This subset is syntactically (and hence semantically) equivalent to the formula specified by the theorem, but is possibly smaller due to the pruning achieved by the IF statement in lines 7–9. A timeline is pruned from the set if it is (semantically) subsumed by any other timeline in the set (one timeline is kept from any semantically equivalent group of timelines, at random). This pruning of timelines is sound, since a timeline is pruned from the output only if it subsumes some other formula in the output—this fact allows an easy argument that the pruned formula is syntactically equivalent to (i.e. mutually syntactically subsumed by) the unpruned formula. Section 5.4.2 traces the computations of this algorithm for an example LGG calculation.

The method does an exponential amount of work even if the result is small (typically because many timelines can be pruned from the output because they subsume what remains). It is still an open question as to whether there is an output-efficient algorithm for computing the syntactic AMA LGG—this problem is in coNP and we conjecture that it is coNP-complete. One route to settling this question is to determine the output complexity of semantic LGG for MA input formulas. We believe that problem also to be coNP-complete, but have not proven this; if that problem is in P, there is an output-efficient method for computing syntactic AMA LGG based on Theorem 25.

A summary of the algorithmic complexity results from this section can be found in Table 3 in the conclusions section of this paper.

```

1: syntactic-LGG( $\{\Psi_1, \Psi_2, \dots, \Psi_m\}$ )
2:   // The  $\Psi_i$  are AMA formulas.
3:   // Outputs a syntactic LGG of the  $\Psi_i$ .
4:    $G = \{\}$ ;
5:   For each  $\langle \Phi_1, \dots, \Phi_m \rangle \in \Psi_1 \times \dots \times \Psi_m$ 
6:     For each  $\Phi$  in semantic-LGG( $\{\Phi_1, \dots, \Phi_m\}$ )
7:       IF ( $\forall \Phi' \in G . \Phi' \not\leq \Phi$ )
8:         THEN  $G' = \{\Phi'' \in G \mid \Phi \leq \Phi''\}$ ;
9:          $G = (G - G') \cup \{\Phi\}$ ;
10:  RETURN ( $\bigwedge G$ )
    
```

Figure 7: Pseudo-code that computes the syntactic AMA LGG of a set of AMA formulas.

5.4 Examples: Least-General Generalization Calculations

Below we work through the details of a semantic and a syntactic LGG calculation. We consider the AMA formulas $\Psi = (A; B) \wedge (B; A)$ and $\Phi = A; B; A$, for which the semantic LGG is $A; B; A$ and the syntactic LGG is $(A; B; \text{true}) \wedge (\text{true}; B; A)$.

5.4.1 SEMANTIC LGG EXAMPLE

The first step in calculating the semantic LGG, according to the algorithm given in Figure 6, is to compute the interdigitation-specializations of the input formulas (i.e., $\text{IS}(\Phi)$ and $\text{IS}(\Psi)$). Trivially, we have that $\text{IS}(\Phi) = \Phi = A; B; A$. To calculate $\text{IS}(\Psi)$, we must consider the possible interdigitations of Ψ , for which there are three,

$$\begin{aligned}
 & \{ \langle A, B \rangle, \langle B, B \rangle, \langle B, A \rangle \} \\
 & \{ \langle A, B \rangle, \langle B, A \rangle \} \\
 & \{ \langle A, B \rangle, \langle A, A \rangle, \langle B, A \rangle \}
 \end{aligned}$$

Each interdigitation leads to the corresponding member of $\text{IS}(\Psi)$ by unioning (conjoining) the states in each tuple, so $\text{IS}(\Psi)$ is

$$\begin{aligned}
 & \{ (A \wedge B); B; (A \wedge B), \\
 & \quad (A \wedge B), \\
 & \quad (A \wedge B); A; (A \wedge B) \}.
 \end{aligned}$$

Lines 5–9 of the semantic LGG algorithm compute the set S , which is equal to the union of the timelines in $\text{IS}(\Psi)$ and $\text{IS}(\Phi)$, with all subsumed timelines removed. For our formulas, we see that each timeline in $\text{IS}(\Psi)$ is subsumed by Φ —thus, we have that $S = \Phi = A; B; A$.

After computing S , the algorithm returns the conjunction of timelines in $\text{IG}(S)$, with redundant timelines removed (i.e., all subsuming timelines are removed). In our case, $\text{IG}(S) = A; B; A$, trivially, as there is only one timeline in S , thus the algorithm correctly computes the semantic LGG of Ψ and Φ to be $A; B; A$.

5.4.2 SYNTACTIC LGG EXAMPLE

The syntactic LGG algorithm, shown in Figure 7, computes a series of semantic LGGs for MA timeline sets, returning the conjunction of the results (after pruning). Line 5 of the algorithm, cycles through timeline tuples from the cross-product of the input AMA formulas. In our case the tuples in $\Phi \times \Psi$ are $T_1 = \langle A; B; A, A; B \rangle$ and $T_2 = \langle A; B; A, B; A \rangle$ —for each tuple, the algorithm computes the semantic LGG of the tuple’s timelines.

The semantic LGG computation for each tuple uses the algorithm given in Figure 6, but the argument is always a set of MA timelines rather than AMA formulas. For this reason, lines 4–9 are superfluous, as for an MA timeline Φ' , $\text{IS}(\Phi') = \Phi'$. In the case of tuple T_1 , lines 4–9 of the algorithm just compute $S = \{A; B; A, A; B\}$. It remains to compute the interdigitation-generalizations of S (i.e., $\text{IG}(S)$), returning the conjunction of those timelines after pruning (lines 10–15 in Figure 6). The set of all interdigitations of S are,

$$\begin{aligned} & \{ \langle A, A \rangle, \langle B, A \rangle, \langle B, B \rangle, \langle B, A \rangle \} \\ & \{ \langle A, A \rangle, \langle B, B \rangle, \langle B, A \rangle \} \\ & \{ \langle A, A \rangle, \langle A, B \rangle, \langle B, B \rangle, \langle B, A \rangle \} \\ & \{ \langle A, A \rangle, \langle A, B \rangle, \langle B, A \rangle \} \\ & \{ \langle A, A \rangle, \langle A, B \rangle, \langle A, A \rangle, \langle B, A \rangle \} \end{aligned}$$

By intersecting states in interdigitation tuples we get $\text{IG}(S)$,

$$\{ A; \mathbf{true}; B; \mathbf{true}, A; B; \mathbf{true}, A; \mathbf{true}; B; \mathbf{true}, A; \mathbf{true}; \mathbf{true}, A; \mathbf{true}; A; \mathbf{true} \}$$

Since the timeline $A; B; \mathbf{true}$ is subsumed by all timelines in $\text{IG}(S)$, all other timelines will be pruned. Thus the semantic LGG algorithm returns $A; B; \mathbf{true}$ as the semantic LGG of the timelines in T_1 .

Next the syntactic LGG algorithm computes the semantic LGG of the timelines in T_2 . Following the same steps as for T_1 , we find that the semantic LGG of the timelines in T_2 is $\mathbf{true}; B; A$. Since $A; B; \mathbf{true}$ and $\mathbf{true}; B; A$ do not subsume one another, the set G computed by lines 5–9 of the syntactic LGG algorithm is equal to $\{ A; B; \mathbf{true}, \mathbf{true}; B; A \}$. Thus, the algorithm computes the syntactic LGG of Φ and Ψ to be $(A; B; \mathbf{true}) \wedge (\mathbf{true}; B; A)$. Note that, in this case, the syntactic LGG is more general than the semantic LGG.

6. Practical Extensions

We have implemented a specific-to-general AMA learning algorithm based on the LGCF and syntactic LGG algorithms presented earlier. This implementation includes three practical extensions. The first extension aims at controlling the exponential complexity by limiting the length of the timelines we consider. The second extension deals with applying our propositional algorithm to relational data, as is necessary for the application domain of visual event recognition. Finally, we show how to gain the practical advantages of negation despite the fact that AMA does not include negation—this turns out to be crucial to achieving good performance in our experiments.

6.1 k -AMA Least-General Generalization

We have already indicated that our syntactic AMA LGG algorithm takes exponential time relative to the lengths of the timelines in the AMA input formulas. This motivates restricting the AMA language to k -AMA in practice, where formulas contain timelines with no more than k states. As k is increased the algorithm is able to output increasingly specific formulas at the cost of an exponential increase in computational time. In the visual-event-recognition experiments shown later, as we increased k , the resulting formulas became overly specific before a computational bottle-neck is reached—i.e., for that application the best values of k were practically computable and the ability to limit k provided a useful language bias.

We use a k -cover operator in order to limit our syntactic LGG algorithm to k -AMA. A k -cover of an AMA formula is a syntactically least general k -AMA formula that syntactically subsumes the input—it is easy to show that a k -cover for a formula can be formed by conjoining all k -MA timelines that syntactically subsume the formula (i.e., that subsume any timeline in the formula). Figure 8 gives pseudo-code for computing the k -cover of an AMA formula. It can be shown that this algorithm correctly computes a k -cover for any input AMA formula. The algorithm calculates the set of least general k -MA timelines that subsume each timeline in the input—the resulting k -MA formulas are conjoined and “redundant” timelines are pruned

```

1: k-cover( $k, \bigwedge_{1 \leq i \leq m} \Phi_i$ )
2:     // The  $\Phi_i$  are AMA formulas.
3:     // Outputs a  $k$ -cover of conjunction of the  $\Phi_i$ .
4:      $G = \{\}$ ;
5:     For  $i = 1$  to  $m$ 
6:         For each  $P = \langle P_1, \dots, P_n \rangle$  in all-values(a-k-partition( $k, \Phi_i$ ))
7:              $\Phi = (\bigcap P_1); \dots; (\bigcap P_n)$ ;
8:             IF  $(\forall \Phi' \in G. \Phi' \not\leq \Phi)$ 
9:                 THEN  $G' = \{\Phi'' \in G \mid \Phi \leq \Phi''\}$ ;
10:                  $G = (G - G') \cup \{\Phi\}$ ;
11:     RETURN  $(\bigwedge G)$ 

12: a-k-partition( $k, s_1; \dots; s_j$ )
13:     //  $s_1; \dots; s_j$  is an MA timeline.
14:     //  $k$  is an positive natural number.
15:     // Non-deterministically partitions  $s_1; \dots; s_j$  into a tuple of  $\leq k$  sets of consecutive states.
16:     IF  $j \leq k$  THEN RETURN  $\langle \{s_1\}, \dots, \{s_j\} \rangle$ ;
17:     IF  $k == 1$  THEN RETURN  $\langle \{s_1, \dots, s_j\} \rangle$ ;
18:      $l = \text{a-member-of}(\{1, 2, \dots, j - k + 1\})$ ; // non-deterministically choose next block size
19:      $P_0 = \{s_1, \dots, s_l\}$ ; // construct next block
20:     RETURN extend-tuple( $P_0, \text{a-k-partition}(k - 1, s_{l+1}; \dots; s_j)$ );
    
```

Figure 8: Pseudo-code for non-deterministically computing a k -cover of an AMA formula, along with a non-deterministic helper function for selecting a $\leq k$ block partition of the states of a timeline.

using a subsumption test. We note that the k -cover of an AMA formula may itself be exponentially larger than that formula; however, in practice, we have found k -covers not to exhibit undue size growth.

Given the k -cover algorithm we restrict our learner to k -AMA as follows: 1) Compute the k -cover for each AMA input formulas. 2) Compute the syntactic AMA LGG of the resulting k -AMA formulas. 3) Return the k -cover of the resulting AMA formula. The primary bottleneck of the original syntactic LGG algorithm is computing the exponentially large set of interdigitation-generalizations—the k -limited algorithm limits this complexity as it only computes interdigitation-generalizations involving k -MA timelines.

6.2 Relational Data

LEONARD produces relational models that involve objects and (force dynamic) relations between those objects. Thus event definitions include variables to allow generalization over objects. For example, a definition for $\text{PICKUP}(x, y, z)$ recognizes both $\text{PICKUP}(\mathbf{hand}, \mathbf{block}, \mathbf{table})$ as well as $\text{PICKUP}(\mathbf{man}, \mathbf{dumbbell}, \mathbf{floor})$. Despite the fact that our k -AMA learning algorithm is propositional, we are still able to use it to learn relational definitions.

We take a straightforward object-correspondence approach to relational learning. We view the models output by LEONARD as containing relations applied to constants. Since we (currently) support only supervised learning, we have a set of distinct training examples for each event type. There is an implicit correspondence between the objects filling the same role across the different training models for a given type. For example, models showing $\text{PICKUP}(\mathbf{hand}, \mathbf{block}, \mathbf{table})$ and $\text{PICKUP}(\mathbf{man}, \mathbf{dumbbell}, \mathbf{floor})$ have implicit

correspondences $\langle \text{hand, man} \rangle$, $\langle \text{block, dumbbell} \rangle$, and $\langle \text{table, floor} \rangle$. We outline two relational learning methods that differ in how much object-correspondence information they require as part of the training data.

6.2.1 COMPLETE OBJECT CORRESPONDENCE

This first approach assumes that a complete object correspondence is given, as input, along with the training examples. Given such information, we can propositionalize the training models by replacing corresponding objects with unique constants. The propositionalized models are then given to our propositional k -AMA learning algorithm which returns a propositional k -AMA formula. We then lift this propositional formula by replacing each constant with a distinct variable. Lavrac et al. (1991) has taken a similar approach.

6.2.2 PARTIAL OBJECT CORRESPONDENCE

The above approach assumes complete object-correspondence information. While it is sometimes possible to provide all correspondences (for example, by color-coding objects that fill identical roles when recording training movies), such information is not always available. When only a partial object correspondence (or even none at all) is available, we can automatically complete the correspondence and apply the above technique.

For the moment, assume that we have an evaluation function that takes two relational models and a candidate object correspondence, as input, and yields an evaluation of correspondence quality. Given a set of training examples with missing object correspondences, we perform a greedy search for the best set of object-correspondence completions over the models. Our method works by storing a set P of propositionalized training examples (initially empty) and a set U of unpropositionalized training examples (initially the entire training set). For the first step, when P is empty, we evaluate all pairs of examples from U , under all possible correspondences, select the pair that yields the highest score, remove the examples involved in that pair from U , propositionalize them according to the best correspondence, and add them to P . For each subsequent step, we use the previously computed values of all pairs of examples, one from U and one from P , under all possible correspondences. We then select the example from U and correspondence that yields the highest average score relative to all models in P —this example is removed from U , propositionalized according to the winning correspondence, and added to P . For a fixed number of objects, the effort expended here is polynomial in the size of the training set; however, if the number of objects b that appear in a training example is allowed to grow, the number of correspondences that must be considered grows as b^b . For this reason, it is important that the events involved manipulate only a modest number of objects.

Our evaluation function is based on the intuition that object roles for visual events (as well as events from other domains) can often be inferred by considering the changes between the initial and final moments of an event. Specifically, given two models and an object correspondence, we first propositionalize the models according to the correspondence. Next, we compute ADD and DELETE lists for each model. The ADD list is the set of propositions that are true at the final moment but not the initial moment. The DELETE list is the set of propositions that are true at the initial moment but not the final moment. (These add and delete lists are motivated by STRIPS action representations. (Fikes & Nilsson, 1971)) Given such ADD_i and DELETE_i lists for models 1 and 2, the evaluation function returns the sum of the cardinalities of $\text{ADD}_1 \cap \text{ADD}_2$ and $\text{DELETE}_1 \cap \text{DELETE}_2$. This heuristic measures the similarity between the ADD and DELETE lists of the two models. The intuition behind this heuristic is similar to the intuition behind the STRIPS action-description language—i.e., that most of the differences between the initial and final moments of an event occurrence are related to the target event, and that event effects can be described by ADD and DELETE lists. We have found that this evaluation function works well in the visual-event domain.

Note, that when full object correspondences are given to the learner (rather than automatically extracted by the learner), the training examples are interpreted as specifying that the target event took place as well as which objects filled the various event roles (e.g., $\text{PICKUP}(a,b,c)$). Rather, when no object correspondences are provided the training examples are interpreted as specifying the existence of a target event occurrence but do not specify which objects fill the roles (i.e., the training example is labeled by PICKUP rather than $\text{PICKUP}(a,b,c)$). Accordingly, the rules learned when no correspondences are provided only allow us

to infer that a target event occurred and not which objects filled the event roles. For example when object correspondences are manually provided the learner might produce the rule,

$$\text{PICKUP}(x, y, z) \triangleq \left[\begin{array}{l} (\text{SUPPORTS}(z, y) \wedge \text{CONTACTS}(z, y)); \\ (\text{SUPPORTS}(x, y) \wedge \text{ATTACHED}(x, y)) \end{array} \right]$$

whereas a learner that automatically extracts the correspondences would instead produce the rule,

$$\text{PICKUP} \triangleq \left[\begin{array}{l} (\text{SUPPORTS}(z, y) \wedge \text{CONTACTS}(z, y)); \\ (\text{SUPPORTS}(x, y) \wedge \text{ATTACHED}(x, y)) \end{array} \right]$$

Its worth noting, however, that upon producing the second rule the availability of a single training example with correspondence information allows the learner to determine the roles of the variables, upon which it can output the first rule. Thus, under the assumption that the learner can reliably extract object correspondences, we need not label all training examples with correspondence information in order to obtain definitions that explicitly recognize object roles.

6.3 Negative Information

The AMA language does not allow negated propositions. Negation, however, is sometimes necessary to adequately define an event type. It turns out that we can easily get the practical advantages of negation without incorporating negation into the AMA language. We do this by adding new propositions to our models that intuitively represent the proposition negations. Assume the training examples contain the propositions $\{p_1, \dots, p_n\}$. We introduce a new set $\{\bar{p}_1, \dots, \bar{p}_n\}$ of propositions and add these into the training models. It is a design choice as to how we assign truth values to these new propositions.

In our experiments, we compare two methods for assigning a truth value to \bar{p}_i . The first method, called *full negation*, assigns true to \bar{p}_i in a model iff p_i is false in the model. The second method, called *boundary negation*, differs from full negation in that it only allows \bar{p}_i to be true in the initial and final moments of a model (and then only if p_i is false). \bar{p}_i must be false at all other times. We have found that boundary negation provides a good trade-off between no negation, which often produces overly general results, and full negation, which often produces overly specific and much more complicated results. Both methods share the property that they produce models where p_i and \bar{p}_i are never simultaneously true. It follows that our learning methods will never produce formulas with states that contain both p_i and \bar{p}_i .

7. Experiments

7.1 Data Set

Our data set contains examples of 7 different event classes: *pick up*, *put down*, *stack*, *unstack*, *move*, *assemble*, and *disassemble*. Each of these involve a hand and two to three blocks. For a detailed description and sample video sequences of these event types, see Siskind (2001). Key frames from sample video sequences of these event classes are shown in figure 9. The results of segmentation, tracking, and model reconstruction are overlaid on the video frames. We recorded 30 movies for each of the 7 event classes resulting in a total of 210 movies comprising 11946 frames. We replaced one *assemble* movie, with a duplicate copy of another because of segmentation and tracking errors.

Some of the event classes are hierarchical in that occurrences of events in one class contain occurrences of events in one or more simpler classes. For example, a movie depicting a $\text{MOVE}(a, b, c, d)$ event (i.e. a moves b from c to d) contains subintervals where $\text{PICKUP}(a, b, c)$ and $\text{PUTDOWN}(a, b, d)$ events occur. In our experiments, when learning the definition of an event class only the movies for that event class are used in training—we do not train on movies for other event classes that may also depict an occurrence of the event class being learned as a sub-event. However, in evaluating the learned definitions, we wish to detect both the events that correspond to an entire movie as well as subevents that correspond to portions of that movie. For example, given a movie depicting $\text{MOVE}(a, b, c, d)$, we wish to detect not only the $\text{MOVE}(a, b, c, d)$

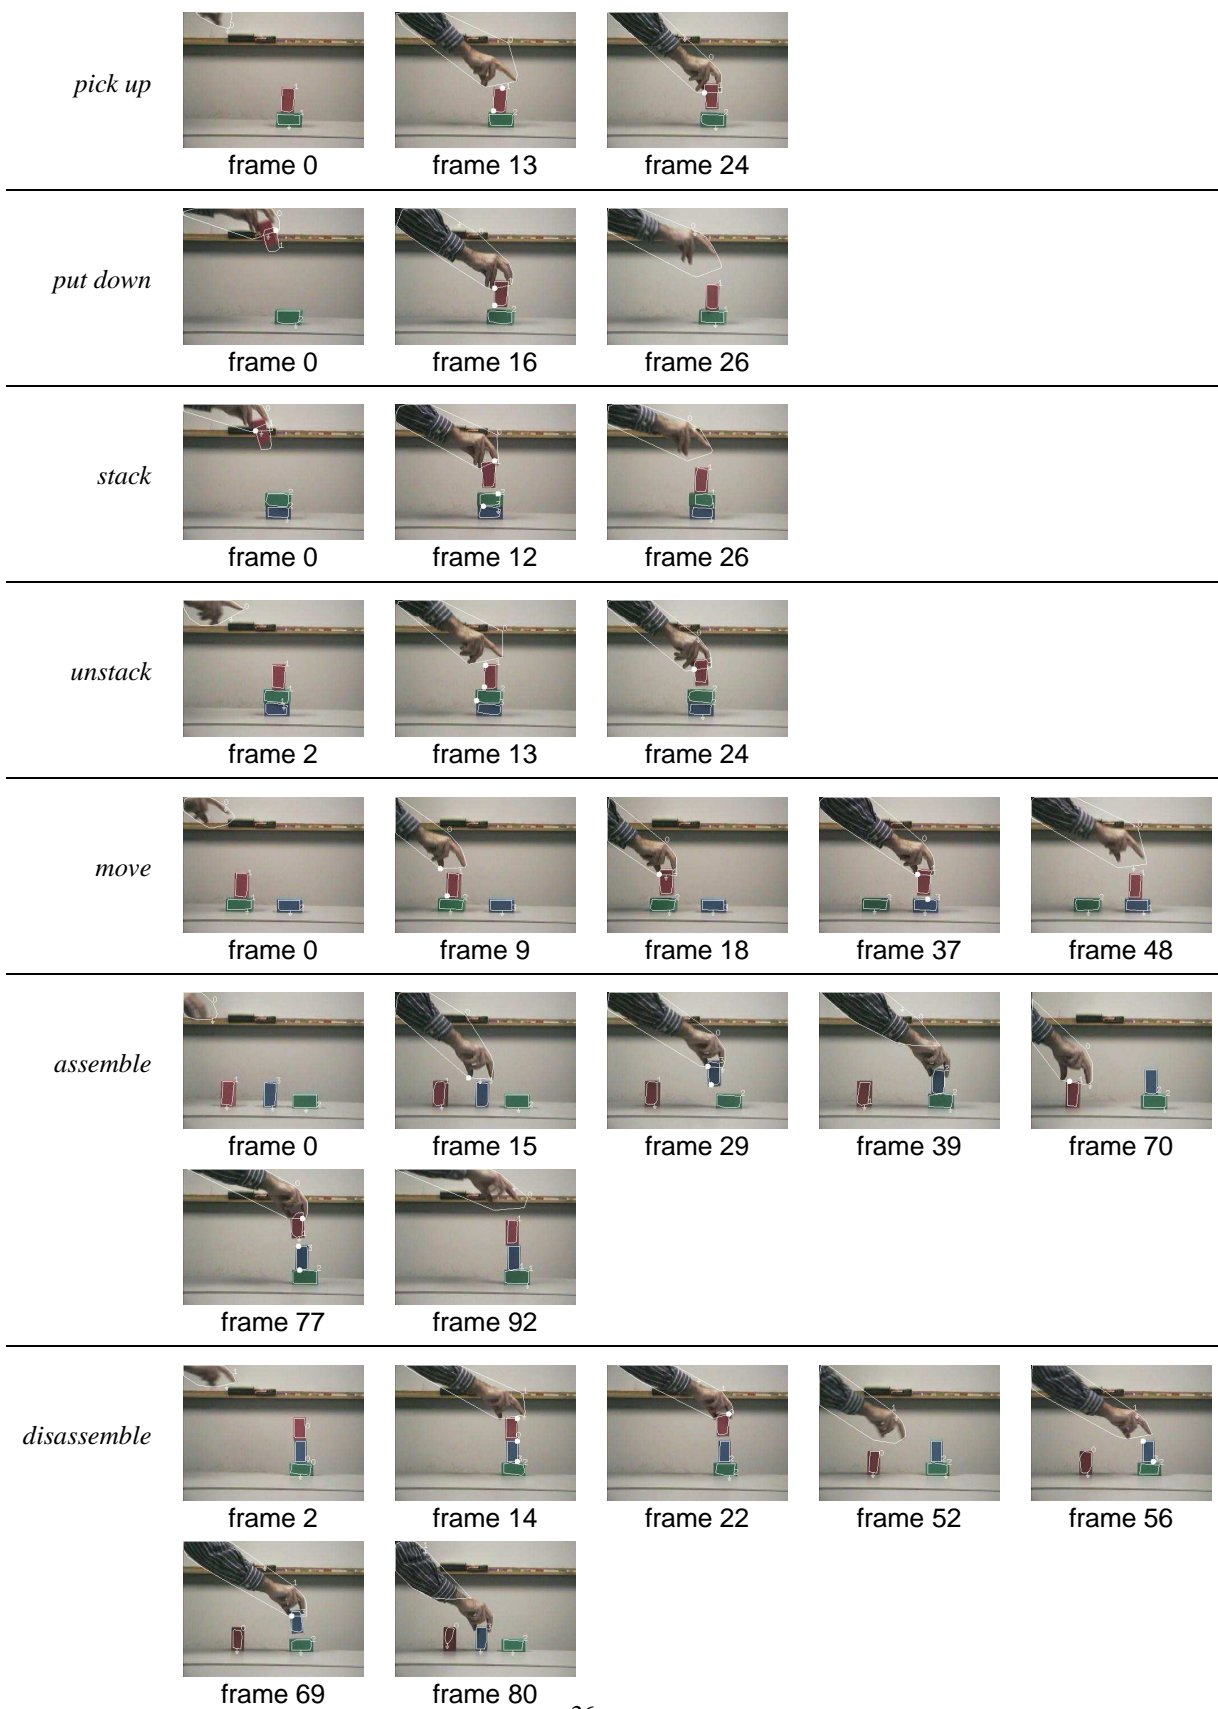


Figure 9: Key frames from sample videos of the seven event types.

event but also the $\text{PICKUP}(a, b, c)$ and $\text{PUTDOWN}(a, b, d)$ subevents as well. For each movie type in our data set, we have a set of *intended* events and subevents that should be detected. If a definition does not detect an intended event, we deem the error a false negative. If a definition detects an unintended event, we deem the error a false positive. For example, if a movie depicts a $\text{MOVE}(a, b, c, d)$ event, the intended events are $\text{MOVE}(a, b, c, d)$, $\text{PICKUP}(a, b, c)$, and $\text{PUTDOWN}(a, b, c)$. If the definition for *pick up* detects the occurrence of $\text{PICKUP}(c, b, a)$ and $\text{PICKUP}(b, a, c)$, but not $\text{PICKUP}(a, b, c)$, it will be charged two false positives as well as one false negative. We evaluate our definitions in terms of false positive and negative rates as describe below.

7.2 Experimental Procedure

For each event type, we evaluate the k -AMA learning algorithm using a leave-one-movie-out cross-validation technique with training-set sampling. The parameters to our learning algorithm are k and the degree D of negative information used: the value of D is either P, for “positive propositions only”, BN, for “boundary negation”, or N, for “full negation”. The parameters to our evaluation procedure include the target event type E and the training-set size N . Given this information, the evaluation proceeds as follows: For each movie M (the held-out movie) from the 210 movies, apply the k -AMA learning algorithm to a randomly drawn training sample of N movies from the 30 movies of event type E (or 29 movies if M is one of the 30). Use LEONARD to detect all occurrences of the learned event definition in M . Based on E and the event type of M , record the number of false positives and false negatives in M , as detected by LEONARD. Let FP and FN be the total number of false positives and false negatives observed over all 210 held-out movies respectively. Repeat the entire process of calculating FP and FN 10 times and record the averages as $\overline{\text{FP}}$ and $\overline{\text{FN}}$.

Since some event types occur more frequently in our data than others (because simpler events occur as subevents of more complex events but not vice versa), we do not report $\overline{\text{FP}}$ and $\overline{\text{FN}}$ directly. Instead, we normalize $\overline{\text{FP}}$ by dividing by the total number of times LEONARD detected the target event (correctly or incorrectly) within all 210 movies and we normalize $\overline{\text{FN}}$ by dividing by the total number of correct occurrences of the target event within all 210 movies (i.e., the human assessment of the number of occurrences of the target event). The normalized value of $\overline{\text{FP}}$ estimates the probability that the target event did not occur given that it was predicted to occur, while the normalized value of $\overline{\text{FN}}$ estimates the probability that the event was not predicted to occur given that it did occur.

7.3 Results

To evaluate our k -AMA learning approach, we ran leave-one-movie-out experiments, as described above, for varying k , D , and N . The 210 example movies were recorded with color-coded objects to provide complete object-correspondence information. We compared our learned event definitions to the performance of two sets of hand-coded definitions. The first set HD_1 of hand-coded definitions appeared in Siskind (2001). In response to subsequent deeper understanding of the behavior of Leonard’s model reconstruction methods, we manually revised these definitions to yield another set HD_2 of hand-coded definitions that gives a significantly better $\overline{\text{FN}}$ performance at some cost in $\overline{\text{FP}}$ performance. Appendix C gives the event definitions in HD_1 and HD_2 along with a set of machine-generated definitions, produced by the k -AMA learning algorithm, given all training data for $k = 30$ and $D = \text{NPN}$.

7.3.1 OBJECT CORRESPONDENCE

To evaluate our algorithm for finding object correspondences, we ignored the correspondence information provided by color coding and applied the algorithm to all training models for each event type. The algorithm selected the correct correspondence for all 210 training models. Thus, for this data set, the learning results when no correspondence information is given will be identical to those where the correspondences are manually provided, except that in the first case the rules will not specify particular object roles (as discussed

k	D		pu	pd	st	un	mo	as	di
2	BN	\overline{FP}	0	0.14	0	0	0	0.75	0
		\overline{FN}	0	0.19	0.12	0.03	0	0	0
3	BN	\overline{FP}	0	0	0	0	0	0	0
		\overline{FN}	0	0.2	0.45	0.10	0.03	0.07	0.10
4	BN	\overline{FP}	0	0	0	0	0	0	0
		\overline{FN}	0	0.2	0.47	0.12	0.03	0.07	0.17
3	P	\overline{FP}	0.42	0.5	0	0.02	0	0	0
		\overline{FN}	0	0.19	0.42	0.11	0.03	0.03	0.10
3	BN	\overline{FP}	0	0	0	0	0	0	0
		\overline{FN}	0	0.2	0.45	0.10	0.03	0.07	0.10
3	N	\overline{FP}	0	0	0	0	0	0	0
		\overline{FN}	0.04	0.39	0.58	0.16	0.13	0.2	0.2
HD ₁		\overline{FP}	0.01	0.01	0	0	0	0	0
		\overline{FN}	0.02	0.22	0.82	0.62	0.03	1.0	0.5
HD ₂		\overline{FP}	0.13	0.11	0	0	0	0	0
		\overline{FN}	0.0	0.19	0.42	0.02	0.0	0.77	0.0

Table 1: \overline{FP} and \overline{FN} for learned definitions, varying both k and D , and for hand-coded definitions.

in section 6.2.2). Since our evaluation procedure uses role information, the rest of our experiments use the manual correspondence information, provided by color-coding, rather than computing it.

7.3.2 VARYING k

The first three rows of table 1 show the \overline{FP} and \overline{FN} values for all 7 event types for $k \in \{2, 3, 4\}$, $N = 29$ (the maximum), and $D = \text{BN}$. Similar trends were found for $D = \text{P}$ and $D = \text{N}$. The general trend is that, as k increases, \overline{FP} decreases or remains the same and \overline{FN} increases or remains the same. Such a trend is a consequence of our k -cover approach. This is because, as k increases, the k -AMA language contains strictly more formulas. Thus for $k_1 > k_2$, the k_1 -cover of a formula will never be more general than the k_2 -cover. This strongly suggests, but does not prove, that \overline{FP} will be non-increasing with k and \overline{FN} will be non-decreasing with k .

Our results show that 2-AMA is overly general for *put down* and *assemble*, i.e. it gives high \overline{FP} . In contrast, 3-AMA achieves $\overline{FP} = 0$ for each event type but pays a penalty in \overline{FN} compared to 2-AMA. Since 3-AMA achieves $\overline{FP} = 0$ there is likely no advantage in moving to k -AMA for $k > 3$ —i.e., the expected result is for \overline{FN} to become larger. This effect is demonstrated for 4-AMA in the table.

7.3.3 VARYING D

Rows four through six of table 1 show \overline{FP} and \overline{FN} for all 7 event types for $D \in \{\text{P}, \text{BN}, \text{N}\}$, $N = 29$, and $k = 3$. Similar trends were observed for other values of k . The general trend is that, as the degree of negative information increases, the learned event definitions become more specific. In other words, \overline{FP} decreases and \overline{FN} increases. This makes sense since, as more negative information is added to the training models, more specific structure can be found in the data and exploited by the k -AMA formulas. We can see that, with $D = \text{P}$, the definitions for *pick up* and *put down* are overly general, as they produce high \overline{FP} . Alternatively,

with $D = N$, the learned definitions are overly specific, giving $\overline{FP} = 0$, at the cost of high \overline{FN} . In these experiments, as well as others, we have found that $D = BN$ yields the best of both worlds: $\overline{FP} = 0$ for all event types and lower \overline{FN} than achieved with $D = N$.

Experiments not shown here have demonstrated that, without negation for *pick up* and *put down*, we can increase k arbitrarily, in an attempt to specialize the learned definitions, and never significantly reduce \overline{FP} . This indicates that negative information plays a particularly important role in constructing definitions for these event types.

7.3.4 COMPARISON TO HAND-CODED DEFINITIONS

The bottom two rows of table 1 show the results for HD_1 and HD_2 . We have not yet attempted to automatically select the parameters for learning (i.e. k and D). Rather, here we focus on comparing the hand-coded definitions to the parameter set that we judged to be best performing across all event types. We believe, however, that these parameters could be selected reliably using cross-validation techniques on a larger data set. In that case, the parameters would be selected on a per-event-type basis and would likely result in an even more favorable comparison to the hand-coded definitions.

The results show that the learned definitions significantly outperform HD_1 on the current data set. The HD_1 definitions were found to produce a large number of false negatives on the current data set. Manual revision of HD_1 yielded HD_2 . Notice that, although HD_2 produces significantly fewer false negatives for all event types, it produces more false positives for *pick up* and *put down*. This is because the hand definitions utilize *pick up* and *put down* as macros for defining the other events.

The performance of the learned definitions is competitive with the performance of HD_2 . The main differences in performance are: (a) for *pick up* and *put down*, the learned and HD_2 definitions achieve nearly the same \overline{FN} but the learned definitions achieve $\overline{FP} = 0$ whereas HD_2 has significant \overline{FP} , (b) for *unstack* and *disassemble*, the learned definitions perform moderately worse than HD_2 with respect to \overline{FN} , and (c) the learned definitions perform significantly better than HD_2 on *assemble* events.

We conjecture that further manual revision could improve HD_2 to perform as well as (and perhaps better than) the learned definitions for every event class. Nonetheless, we view this experiment as promising, as it demonstrates that our learning technique is able to compete with, and sometimes outperform, significant hand-coding efforts by a domain expert.

7.3.5 VARYING N

It is of practical interest to know how training set size affects our algorithm’s performance. For this application, it is important that our method work well with fairly small data sets, as it can be tedious to collect event data. Table 2 shows the \overline{FN} of our learning algorithm for each event type, as N is reduced from 29 to 5. For these experiments, we used $k = 3$ and $D = BN$. Note that $\overline{FP} = 0$ for all event types and all N and hence is not shown. We expect \overline{FN} to increase as N is decreased, since, with specific-to-general learning, more data yields more-general definitions. Generally, \overline{FN} is flat for $N > 20$, increases slowly for $10 < N < 20$, and increases abruptly for $5 < N < 10$. We also see that, for several event types, \overline{FN} decreases slowly, as N is increased from 20 to 29. This indicates that a larger data set might yield improved results for those event types.

8. Related Work

Here we discuss two bodies of related work. First, we present previous work in visual event recognition and how it relates to our experiments here. Second, we discuss previous approaches to learning temporal patterns from positive data.

N	pu	pd	st	un	mo	as	di
29	0.0	0.20	0.45	0.10	0.03	0.07	0.10
25	0.0	0.20	0.47	0.16	0.05	0.09	0.10
20	0.01	0.21	0.50	0.17	0.08	0.12	0.12
15	0.01	0.22	0.53	0.26	0.14	0.20	0.16
10	0.07	0.27	0.60	0.36	0.23	0.32	0.26
5	0.22	0.43	0.77	0.54	0.35	0.57	0.43

Table 2: \overline{FN} for $k = 3$, $D = \text{BN}$, and various values of N .

8.1 Visual Event Recognition

Prior work has investigated various subsets of the pieces of learning and using temporal, relational, and force-dynamic representations for recognizing events in video. But none, to date, combine all the pieces together. The following is a representative list and not meant to be comprehensive. Borchardt (1985) presents temporal, relational, force-dynamic event definitions but these definitions are neither learned nor applied to video. Regier (1992) presents techniques for learning temporal event definitions but the learned definitions are neither relational, force dynamic, nor applied to video. Yamoto, Ohya, and Ishii (1992), Brand and Essa (1995), Brand, Oliver, and Pentland (1997), and Bobick and Ivanov (1998) present techniques for learning temporal event definitions from video but the learned definitions are neither relational nor force dynamic. Pinhanez and Bobick (1995) and Brand (1997a) present temporal, relational event definitions that recognize events in video but these definitions are neither learned nor force dynamic. Brand (1997b) and Mann and Jepson (1998) present techniques for analyzing force dynamics in video but neither formulate event definitions nor apply these techniques to recognizing events or learning event definitions.

8.2 Learning Temporal Patterns

We divide this body of work into three main categories: temporal data mining, inductive logic programming, and finite-state machine induction.

Temporal Data Mining. The sequence-mining literature contains many general-to-specific (“levelwise”) algorithms for finding frequent sequences (Agrawal & Srikant, 1995; Mannila, Toivonen, & Verkamo, 1995; Kam & Fu, 2000; Cohen, 2001; Hoppner, 2001). Here we explore a specific-to-general approach. In this previous work, researchers have studied the problem of mining temporal patterns using languages that are interpreted as placing constraints on partially or totally ordered sets of time-points, e.g., sequential patterns (Agrawal & Srikant, 1995) and episodes (Mannila et al., 1995). These languages place constraints on time-points rather than time-intervals as in our work here. More recently there has been work on mining temporal patterns using interval-based pattern languages (Kam & Fu, 2000; Cohen, 2001; Hoppner, 2001).

Though the languages and learning frameworks vary among these approaches they share two central features which distinguish them from our approach. First, they all typically have the goal of finding all “frequent” patterns (formulas) within a temporal data set—our approach is focused on finding patterns with a frequency of one (covering all positive examples). Our first learning application of visual event recognition has not yet required us to find patterns with frequency less than one; however, there are a number of ways in which we can extend our method in that direction when it becomes necessary (e.g., to deal with noisy training data). Second, these approaches all use standard general-to-specific “level-wise” search techniques, whereas we chose to take a specific-to-general approach. One direction for future work is to develop a general-to-specific “level-wise” algorithm for finding frequent MA formulas and to compare it with our specific-to-general approach. Another direction is to design a “level-wise” version of our specific-to-general algorithm—where for example, the results obtained for the k -AMA LGG can be used to more efficiently calculate the $(k+1)$ -AMA LGG. Whereas a “level-wise” approach is conceptually straightforward in a general-to-specific

framework it is not so clear in the specific-to-general case. We are not familiar with other temporal data-mining systems that take a specific-to-general approach.

First-Order Learning In Section 4.3, we pointed out difficulties in using existing first-order clausal generalization techniques for learning AMA formulas. In spite of these difficulties, it is still possible to represent temporal events in first-order logic (either with or without capturing the AMA semantics precisely) and to apply general-purpose relational learning techniques, e.g., inductive logic programming (ILP) (Muggleton & De Raedt, 1994). Most ILP systems require both positive and negative training examples and hence are not suitable for our current positive-only framework—exceptions include Golem (Muggleton & Feng, 1992), Claudien (De Raedt & Dehaspe, 1997), and Progol (Muggleton, 1995), among others. While we have not performed a full evaluation of these systems, our early experiments in the visual event recognition domain confirmed our belief that horn clauses, lacking special handling of time, give a poor inductive bias. In particular, many of the learned clauses find patterns that simply “do not make sense” from a temporal perspective, and in turn generalize poorly. We believe a reasonable alternative to our approach may be to incorporate syntactic biases into ILP systems as done, for example, in (Cohen, 1994; Dehaspe & De Raedt, 1996; Klingspor, Morik, & Rieger, 1996). In this work, however, we chose to work directly in a temporal logic representation.

Finite-State Machines Finally, we note there has been much theoretical and empirical research into learning finite-state machines (FSMs) (Angluin, 1987; Lang, Pearlmuter, & Price, 1998). We can view FSMs as describing properties of strings (symbol sequences). In our case, however, we are interested in describing sequences of propositional models rather than just sequences of symbols. This suggests learning a type of ‘factored’ FSM where the arcs are labeled by sets of propositions rather than by single symbols. Factored FSM may be a natural direction in which to extend the expressiveness of our current language (for example by allowing repetition). We are not aware of work concerned with learning ‘factored’ FSMs; however, it is likely that inspiration can be drawn from symbol-based FSM learning algorithms.

9. Conclusion

We have presented a simple logic for representing temporal events called AMA and have shown theoretical and empirical results for learning AMA formulas. Empirically, we’ve given the first system for learning temporal, relational, force-dynamic event definitions from positive-only input and we have applied that system to learn such definitions from real video inputs. The resulting performance matches that of event definitions that are hand-coded with substantial effort by human domain experts. On the theoretical side, Table 3 summarizes the upper and lower bounds we have shown for the subsumption and generalization problems associated with this logic. In each case, we have provided a provably correct algorithm matching the upper bound shown. The table also shows the worst-case size that the smallest LGG could possibly take relative to the input size, for both AMA and MA inputs. The key results in this table are the polynomial-time MA subsumption and AMA syntactic subsumption, the coNP lower bound for AMA subsumption, the exponential size of LGGs in the worst case, and the apparently lower complexity of syntactic AMA LGG versus semantic LGG. We described how to build a learner based on these results and applied it to the visual-event learning domain. To date, however, the definitions we learn are neither cross-modal nor perspicuous. And while the performance of the learned definitions matches that of hand-coded ones, we wish to surpass hand coding. In the future, we intend to address cross-modality by applying our learning technique to the planning domain. We also believe that addressing perspicuity will lead to improved performance.

References

- Agrawal, R., & Srikant, R. (1995). Mining sequential patterns. In *International Conference on Data Engineering*, pp. 3–14.
- Allen, J. F. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11), 832–843.

Inputs	Subsumption		Semantic AMA LGG			Syntactic AMA LGG		
	Semantic	Syntactic	Lower	Upper	Size	Lower	Upper	Size
MA	P	P	P	coNP	EXP	P	coNP	EXP
AMA	coNP-complete	P	coNP	NEXP	2-EXP?	P	coNP	EXP

Table 3: Complexity Results Summary. The LGG complexities are relative to *input plus output* size. The size column reports the worst-case smallest correct output size. The “?” indicates a conjecture.

- Angluin, D. (1987). Learning regular sets from queries and counterexamples. *Information and Computation*, 75, 87–106.
- Bacchus, F., & Kabanza, F. (2000). Using temporal logics to express search control knowledge for planning. *Artificial Intelligence*, 16, 123–191.
- Bobick, A. F., & Ivanov, Y. A. (1998). Action recognition using probabilistic parsing. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 196–202, Santa Barbara, CA.
- Borchardt, G. C. (1985). Event calculus. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pp. 524–527, Los Angeles, CA.
- Brand, M. (1997a). The inverse hollywood problem: From video to scripts and storyboards via causal analysis. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pp. 132–137, Providence, RI.
- Brand, M. (1997b). Physics-based visual understanding. *Computer Vision and Image Understanding*, 65(2), 192–205.
- Brand, M., & Essa, I. (1995). Causal analysis for visual gesture understanding. In *Proceedings of AAAI Fall Symposium on Computational Models for Integrating Language and Vision*.
- Brand, M., Oliver, N., & Pentland, A. (1997). Coupled hidden Markov models for complex action recognition. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.
- Cohen, P. (2001). Fluent learning: Elucidation the structure of episodes. In *Symposium on Intelligent Data Analysis*.
- Cohen, W. (1994). Grammatically biased learning: Learning logic programs using an explicit antecedent description lanugage. *Artificial Intelligence*, 68, 303–366.
- Cohen, W., & Hirsh, H. (1994). Learnability of the classic description logic: Theoretical and experimental results. In *4th International Knowledge Representation and Reasoning*, pp. 121–133.
- De Raedt, L., & Dehaspe, L. (1997). Clausal discovery. *Machine Learning*, 26, 99–146.
- Dehaspe, L., & De Raedt, L. (1996). Dlab: A declarative language bias formalism. In *International Symposium on Methodologies for Intelligent Systems*, pp. 613–622.
- Fikes, R., & Nilsson, N. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3/4).
- Hoppner, F. (2001). Discovery of temporal patterns—learning rules about the qualitative behaviour of time series. In *5th European Principles and Practice of Knowledge Discovery in Databases*.
- Kam, P., & Fu, A. (2000). Discovering temporal patterns for interval-based events. In *International Conference on Data Warehousing and Knowledge Discovery*.
- Klingspor, V., Morik, K., & Rieger, A. D. (1996). Learning concepts from sensor data of a mobile robot. *Artificial Intelligence*, 23(2/3), 305–332.

- Lang, K., Pearlmutter, B., & Price, R. (1998). Results of the abbadingo one dfa learning competition and a new evidence-driven state merging algorithm. In *International Colloquium on Grammatical Inference*.
- Lavrac, N., Dzeroski, S., & Grobelnik, M. (1991). Learning nonrecursive definitions of relations with LINUS. In *Proceedings of the Fifth European Working Session on Learning*, pp. 265–288.
- Mann, R., & Jepson, A. D. (1998). Toward the computational perception of action. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 794–799, Santa Barbara, CA.
- Mannila, H., Toivonen, H., & Verkamo, A. I. (1995). Discovery of frequent episodes in sequences. In *International Conference on Data Mining and Knowledge Discovery*.
- Mitchell, T. (1982). Generalization as search. *Artificial Intelligence*, 18(2), 517–42.
- Muggleton, S. (1995). Inverting entailment and Progol. In *Machine Intelligence*, Vol. 14, pp. 133–188. Oxford University Press.
- Muggleton, S., & Feng, C. (1992). Efficient induction of logic programs. In Muggleton, S. (Ed.), *Inductive Logic Programming*, pp. 281–298. Academic Press.
- Muggleton, S., & De Raedt, L. (1994). Inductive logic programming: Theory and methods. *Journal of Logic Programming*, 19/20, 629–679.
- Pinhanez, C., & Bobick, A. (1995). Scripts in machine understanding of image sequences. In *AAAI Fall Symposium Series on Computational Models for Integrating Language and Vision*.
- Plotkin, G. D. (1971). *Automatic Methods of Inductive Inference*. Ph.D. thesis, Edinburgh University.
- Regier, T. P. (1992). *The Acquisition of Lexical Semantics for Spatial Terms: A Connectionist Model of Perceptual Categorization*. Ph.D. thesis, University of California at Berkeley.
- Roth, D., & Yih, W. (2001). Relational learning via propositional algorithms: An information extraction case study. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*.
- Siskind, J. M. (2001). Grounding the lexical semantics of verbs in visual perception using force dynamics and event logic. *Journal of Artificial Intelligence Research*, 15, 31–90.
- Siskind, J. M., & Morris, Q. (1996). A maximum-likelihood approach to visual event classification. In *Proceedings of the Fourth European Conference on Computer Vision*, pp. 347–360, Cambridge, UK. Springer-Verlag.
- Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, 12, 49–100.
- Yamato, J., Ohya, J., & Ishii, K. (1992). Recognizing human action in time-sequential images using hidden Markov model. In *Proceedings of the 1992 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 379–385. IEEE Press.

Appendix A. Internal Positive Event Logic

Here we give the syntax and semantics for an event logic called *Internal Positive Event Logic (IPEL)*. This logic is used in the main text only to motivate our choice of a small subset of this logic, AMA, by showing, in Proposition 4, that AMA can define any set of models that IPEL can define.

An event type (i.e., set of models) is said to be *internal* if whenever it contains any model $\mathcal{M} = \langle M, I \rangle$, it also contains any model that agrees with \mathcal{M} on truth assignments $M[i]$ where $i \in I$. Full event logic allows the definition of non-internal events, for example: the formula $\Psi = \diamond_{<} P$ is satisfied by $\langle M, I \rangle$ when there is some interval I' entirely preceding I such that P is satisfied by $\langle M, I' \rangle$, thus Ψ is not internal. The applications we are considering do not appear to require non-internal events, thus we currently only consider events that are internal.

Call an event type *positive* if it contains the model $\mathcal{M} = \langle M, [1, 1] \rangle$ where $M(1)$ is the truth assignment assigning all propositions the value “true”. A positive event type cannot require any proposition to be false at any point in time.

I_1	Relation	I_2	English	Definition	Inverse
$[m_1, m_2]$	s	$[n_1, n_2]$	starts	$m_1 = n_1$ and $m_2 \leq n_2$	si
$[m_1, m_2]$	f	$[n_1, n_2]$	finishes	$m_1 \leq n_1$ and $m_2 = n_2$	fi
$[m_1, m_2]$	d	$[n_1, n_2]$	during	$m_1 \geq n_1$ and $m_2 \leq n_2$	di
$[m_1, m_2]$	b	$[n_1, n_2]$	before	$m_2 \leq n_1$	bi
$[m_1, m_2]$	m	$[n_1, n_2]$	meets	$m_2 = n_1$ or $m_2 + 1 = n_1$	mi
$[m_1, m_2]$	o	$[n_1, n_2]$	overlaps	$m_1 \leq n_1 \leq m_2 \leq n_2$	oi
$[m_1, m_2]$	=	$[n_1, n_2]$	equals	$m_1 = n_1$ and $m_2 = n_2$	=

Table 4: The Thirteen Allen Relations (adapted to our semantics).

IPEL is a fragment of full propositional event logic that can only describe positive internal events. We conjecture, but have not yet proven, that all positive internal events representable in the full event logic of Siskind (2001) can be represented by some IPEL formula. Formally, the syntax of IPEL formulas is given by

$$E ::= \mathbf{true} \mid \mathit{prop} \mid E_1 \vee E_2 \mid \diamond_{R'} E_1 \mid E_1 \wedge_R E_2,$$

where the E_i are IPEL formulas, prop is a primitive proposition (sometimes called primitive event-type), R is a subset of the thirteen Allen interval relations $\{s, f, d, b, m, o, =, si, fi, di, bi, ai, oi\}$ (Allen, 1983), and R' is a subset of the restricted set of Allen relations $\{s, f, d, =\}$ (the semantics for each Allen relation is given below). The difference between IPEL syntax and that of full propositional event logic is that event logic allows for a negation operator, and that, in full event logic, R' can be any subset of all thirteen Allen relations. The operators \wedge and $;$ used to define AMA formulas are merely abbreviations for the IPEL operators $\wedge_{\{=\}}$ and $\wedge_{\{m\}}$ respectively, so AMA is a subset of IPEL (though a distinguished subset as indicated by Proposition 4).

Each of the thirteen interval Allen relations are binary relations on the set of closed natural-number intervals. Table 4 gives the definitions of these relations, defining $[m_1, m_2] r [n_1, n_2]$ for each Allen relation r . Satisfiability for IPEL formulas can now be defined as follows,

- \mathbf{true} is satisfied by every model.
- prop is satisfied by model $\langle M, I \rangle$ iff $M[x]$ assigns prop true for every $x \in I$.
- $E_1 \vee E_2$ is satisfied by a model \mathcal{M} iff \mathcal{M} satisfies E_1 or \mathcal{M} satisfies E_2 .
- $\diamond_{R'} E$ is satisfied by model $\langle M, I \rangle$ iff for some $r \in R'$ there is an interval I' such that $I' r I$ and $\langle M, I' \rangle$ satisfies E .
- $E_1 \wedge_R E_2$ is satisfied by model $\langle M, I \rangle$ iff for some $r \in R$ there exists intervals I_1 and I_2 such that $I_1 r I_2$, $\mathit{SPAN}(I_1, I_2) = I$ and both $\langle M, I_1 \rangle$ satisfies E_1 and $\langle M, I_2 \rangle$ satisfies E_2 .

where prop is a primitive proposition, E and E_i are IPEL formulas, R is a set of Allen relations, and $\mathit{SPAN}(I_1, I_2)$ is the minimal interval that contains both I_1 and I_2 . From this definition it is easy to show by induction on the number of operators and connectives in a formula that all IPEL formulas define internal events. One can also verify that the definition of satisfiability given earlier for AMA formulas corresponds to the one we give here.

Appendix B. Omitted Proofs

Lemma 1 *For any MA timeline Φ and any model \mathcal{M} , if \mathcal{M} satisfies Φ then there is a witnessing interdigitation for $\text{MAP}(\mathcal{M}) \leq \Phi$.*

Proof: Assume that $\mathcal{M} = \langle M, I \rangle$ satisfies the MA timeline $\Phi = s_1; \dots; s_n$, and let $\Phi' = \text{MAP}(\mathcal{M})$. It is straightforward to argue, by induction on the length of Φ , that there exists a mapping V' from states of Φ to sub-intervals of I , such that

- for any $i \in V'(s)$, $M[i]$ satisfies s ,
- $V'(s_1)$ includes the initial timepoint of I ,
- $V'(s_n)$ includes the final timepoint of I , and
- for any $i \in [1, n - 1]$, we have $V'(s_i)$ meets $V'(s_{i+1})$ (see Table 4).

Let V be the relation between states $s \in \Phi$ and members $i \in I$ that is true when $i \in V'(s)$. Note that the conditions on V' ensure that every $s \in \Phi$ and every $i \in I$ appear in some tuple in V (not necessarily together). Below we use V to construct a witnessing interdigitation W .

Let R be the total, one-to-one, onto function from time-points in I to corresponding states in Φ' , noting that Φ' has one state for each time-point in I , as $\Phi' = \text{MAP}(\langle M, I \rangle)$. Note that R preserves ordering in that, when $i \leq j$, $R(i)$ is no later than $R(j)$ in Φ' . Let W be the composition $V \circ R$ of the relations V and R .

We show that W is an interdigitation. We first show that each state from Φ or Φ' appears in a tuple in W , so W is piecewise total. States from Φ must appear, trivially, because each appears in a tuple of V , and R is total. States from Φ' appear because each $i \in I$ appears in a tuple of V , and R is onto the states of Φ' .

It now suffices to show that for any states s before t from Φ , $W(s, s')$ and $W(t, t')$ implies that t is no later than t' in Φ' , so that W is simultaneously consistent. The conditions defining V' above imply that every number in $i \in V(s)$ is less than or equal to every $j \in V(t)$. The order-preservation property of R , noted above, then implies that every state $s' \in V \circ R(s)$ is no later than any state $t' \in V \circ R(t)$ in Φ' , as desired. So W is an interdigitation.

We now argue that W witnesses $\Phi' \leq \Phi$. Consider $s \in \Phi$ and $t \in \Phi'$ such that $W(s, t)$. By the construction of W , there must be $i \in V'(s)$ for which t is the i 'th state of Φ' . Since $\Phi' = \text{MAP}(\mathcal{M})$, it follows that t is the set of true propositions in $M[i]$. Since $i \in V'(s)$, we know that $M[i]$ satisfies s . It follows that $s \subseteq t$, and so $t \leq s$. \square

Lemma 3 *For any $E \in \text{IPEL}$, if model \mathcal{M} embeds any model that satisfies E then \mathcal{M} satisfies E .*

Proof: Consider the models $\mathcal{M} = \langle M, I \rangle$ and $\mathcal{M}' = \langle M', I' \rangle$ such that \mathcal{M} embeds \mathcal{M}' , let $\Phi = \text{MAP}(\mathcal{M})$ and $\Phi' = \text{MAP}(\mathcal{M}')$. Assume that $E \in \text{IPEL}$ is satisfied by \mathcal{M}' , we will show that E is also satisfied by \mathcal{M} .

We know from the definition of embedding that $\Phi \leq \Phi'$ and thus there is a witnessing interdigitation W for $\Phi \leq \Phi'$ by Proposition 2. We know there is a one-to-one correspondence between numbers in I (I') and states of Φ (Φ') and denote the state in Φ (Φ') corresponding to $i \in I$ ($i' \in I'$) as s_i ($t_{i'}$). This correspondence allows us to naturally interpret W as a mapping V from subsets of I' to subsets of I as follows: for $I'_1 \subseteq I'$, $V(I'_1)$ equals the set of all $i \in I$ such that for some $i' \in I'_1$, s_i co-occurs with $t_{i'}$ in W . We will use the following properties of V ,

1. If I'_1 is a sub-interval of I' , then $V(I'_1)$ is a sub-interval of I .
2. If I'_1 is a sub-interval of I' , then $\langle M, V(I'_1) \rangle$ embeds $\langle M', I'_1 \rangle$.
3. If I'_1 and I'_2 are sub-intervals of I' , and r is an Allen relation, then $I'_1 r I'_2$ iff $V(I'_1) r V(I'_2)$.
4. If I'_1 and I'_2 are sub-intervals of I' , then $V(\text{SPAN}(I'_1, I'_2)) = \text{SPAN}(V(I'_1), V(I'_2))$.

5. $V(I') = I$.

We sketch the proofs of these properties. 1) Use induction on the length of I'_1 , with the definition of interdigitation. 2) Since $V(I'_1)$ is an interval, $\text{MAP}(\langle M, V(I'_1) \rangle)$ is well defined. $\text{MAP}(\langle M, V(I'_1) \rangle) \leq \text{MAP}(\langle M', I'_1 \rangle)$ follows from the assumption that \mathcal{M} embeds \mathcal{M}' . 3) From Appendix A, we see that all Allen relations are defined in terms of the \leq relation on the natural number endpoints of the intervals. We can show that V preserves \leq (but not $<$) on singleton sets (i.e., every member of $V(\{i'\})$ is \leq every member of $V(\{j'\})$ when $i' \leq j'$) and that V commutes with set union. It follows that V preserves the Allen interval relations. 4) Use the fact that V preserves \leq in the sense just argued, along with the fact that $\text{SPAN}(I'_1, I'_2)$ depends only on the minimum and maximum numbers in I'_1 and I'_2 . 5) Follows from the definition of interdigitation and the construction of V .

We now use induction on the number of operators and connectives in E to prove that, if \mathcal{M}' satisfies E , then so must \mathcal{M} . The base case is when $E = \text{prop}$, where prop is a primitive proposition, or **true**. Since \mathcal{M}' satisfies E , we know that prop is true in all $M'[x']$ for $x' \in I'$. Since W witnesses $\Phi \leq \Phi'$, we know that, if prop is true in $M'[x]$, then prop is true in all $M[x]$, where $x \in V(x')$. Therefore, since $V(I') = I$, prop is true for all $M'[x]$, where $x \in I$, hence \mathcal{M}' satisfies E .

For the inductive case, assume that the claim holds for IPEL formulas with fewer than N operators and connectives—let E_1, E_2 be two such formulas. When $E = E_1 \vee E_2$, the claim trivially holds. When $E = \diamond_R E_1$, R must be a subset of the set of relations $\{s, f, d, =\}$. Notice that E can be written as a disjunction of $\diamond_r E_1$ formulas, where r is a single Allen relation from R . Thus, it suffices to handle the case where R is a single Allen relation. Suppose $E = \diamond_{\{s\}} E_1$. Since \mathcal{M}' satisfies E , there must be a sub-interval I'_1 of I' such that $I'_1 s I'$ and $\langle M', I'_1 \rangle$ satisfies E_1 . Let $I_1 = V(I'_1)$, we know from the properties of V that $V(I') = I$, and, hence, that $I_1 s I$. Furthermore, we know that $\langle M, I_1 \rangle$ embeds $\langle M', I'_1 \rangle$, and, thus, by the inductive hypothesis, $\langle M, I_1 \rangle$ satisfies E_1 . Combining these facts, we get that E is satisfied by \mathcal{M} . Similar arguments hold for the remaining three Allen relations. Finally, consider the case when $E = E_1 \wedge_R E_2$, where R can be any set of Allen relations. Again, it suffices to handle the case when R is a single Allen relation r . Since \mathcal{M}' satisfies $E = E_1 \wedge_r E_2$, we know that there are sub-intervals I'_1 and I'_2 of I' such that $\text{SPAN}(I'_1, I'_2) = I'$, $I'_1 r I'_2$, $\langle M', I'_1 \rangle$ satisfies E_1 , and $\langle M', I'_2 \rangle$ satisfies E_2 . From these facts, and the properties of V , it is easy to verify that \mathcal{M} satisfies E . \square

Lemma 5 *Given an MA formula Φ that subsumes each member of a set Σ of MA formulas, Φ also subsumes some member Φ' of $\text{IG}(\Sigma)$. Dually, when Φ is subsumed by each member of Σ , we have that Φ is also subsumed by some member Φ' of $\text{IS}(\Sigma)$. In each case, the length of Φ' can be bounded by the size of Σ .*

Proof: We prove the result for $\text{IG}(\Sigma)$ —the proof for $\text{IS}(\Sigma)$ follows similar lines. Let $\Sigma = \{\Phi_1, \dots, \Phi_n\}$, $\Phi = s_1; \dots; s_m$, and assume that for each $1 \leq i \leq n$, $\Phi_i \leq \Phi$. From Proposition 2, for each i , there is a witnessing interdigitation W_i for $\Phi_i \leq \Phi$. We will combine the W_i into an interdigitation of Σ , and show that the corresponding member of $\text{IG}(\Sigma)$ is subsumed by Φ .

To construct an interdigitation of Σ , first notice that, for each s_j , each W_i specifies a set of states (possibly a single state but at least one) from Φ_i that all co-occur with s_j . Furthermore, since W_i is an interdigitation, it is easy to show that this set of states corresponds to a consecutive sub-sequence of states from Φ_i —let $\Phi_{j,i}$ be the MA timeline corresponding to this subsequence. Now let $\Sigma_j = \{\Phi_{j,i} \mid 1 \leq i \leq n\}$, and α_j be any interdigitation of Σ_j . We now take I to be the union of all α_j , for $1 \leq j \leq m$.

We show that I is an interdigitation of Σ . Since each state s appearing in Σ must co-occur with at least one state s_j in Φ in at least one W_i , s will be in at least one tuple of α_j , and, hence, be in some tuple of I —so I is piecewise total.

Now, define the restriction $I^{i,j}$ of I to components i and j , with $i < j$, to be the relation given by taking the set of all pairs formed by shortening tuples of I by omitting all components except the i 'th and the j 'th. Likewise define $\alpha_k^{i,j}$ for each k . To show I is an interdigitation, it now suffices to show that each $I^{i,j}$ is simultaneously consistent. Consider states s_i and s_j from timelines Φ_i and Φ_j , respectively, such that $I^{i,j}(s_i, s_j)$. Suppose that t_i occurs after s_i in Φ_i , and for some $t_j \in \Phi_j$, $I^{i,j}(t_i, t_j)$ holds. It suffices to show

that s_j is no later than t_j in Φ_j . Since $I^{i,j}(s_i, s_j)$ and $I^{i,j}(t_i, t_j)$, we must have $\alpha_k^{i,j}(s_i, s_j)$ and $\alpha_{k'}^{i,j}(t_i, t_j)$, respectively, for some k and k' . We know $k \leq k'$ because s_i is before t_i in Φ_i and W_i is simultaneously consistent. If $k = k'$, then s_j is no later than t_j in Φ_j , because α_k must be simultaneously consistent, being an interdigitation. Otherwise, $k < k'$. Then s_j is no later than t_j in Φ_j , as desired, because W_j is simultaneously consistent. So I is simultaneously consistent, and an interdigitation of Σ .

Let Φ' be the member of $\text{IG}(\Sigma)$ corresponding to I . We now show that $\Phi' \leq \Phi$. We know that each state $s' \in \Phi'$ is the intersection of the states in a tuple of some α_j —we say that s' derives from α_j . Consider the interdigitation I' between Φ and Φ' , where $I(s_j, s')$, for $s_j \in \Phi$ and $s' \in \Phi'$, if and only if s' derives from I_j . I' is piecewise total, as every tuple of I derives from some α_j , and no α_j is empty. I' is simultaneously consistent because tuples of I deriving from later α_k must be later in the lexicographic ordering of I , given the simultaneous consistency of the W_k interdigitations used to construct each α_j . Finally, we know that s_j subsumes (i.e., is a subset of) each state in each tuple of α_j , because each W_k is a witnessing interdigitation to $\Phi_k \leq \Phi$, and, hence, subsumes (is a subset of) the intersection of those states. Therefore, if $s_j \in \Phi$ co-occurs with $s' \in \Phi'$ in I' we have that $s' \leq s_j$. Thus, I' is a witnessing interdigitation for $\Phi' \leq \Phi$, and by Proposition 2 we have $\Phi' \leq \Phi$.

The size bound on Φ' follows, since, as pointed out in the main text, the size of any member of $\text{IG}(\Sigma)$ is upper-bounded by the number of states in Σ . \square

Lemma 8 *Given MA timelines $\Phi_1 = s_1; \dots; s_m$ and $\Phi_2 = t_1; \dots; t_n$, there is a witnessing interdigitation for $\Phi_1 \leq \Phi_2$ iff there is a path in the subsumption graph $SG(\Phi_1, \Phi_2)$ from $v_{1,1}$ to $v_{m,n}$.*

Proof: The subsumption graph $SG(\Phi_1, \Phi_2)$ is equal to $\langle V, E \rangle$ with $V = \{v_{i,j} \mid 1 \leq i \leq m, 1 \leq j \leq n\}$ and $E = \{\langle v_{i,j}, v_{i',j'} \rangle \mid s_i \leq t_j, s_{i'} \leq t_{j'}, i \leq i' \leq i+1, j \leq j' \leq j+1\}$. Note that there is a correspondence between vertices and state tuples—with vertex $v_{i,j}$ corresponding to $\langle s_i, t_j \rangle$.

For the forward direction, assume that W is a witnessing interdigitation for $\Phi_1 \leq \Phi_2$. We know that, if the states s_i and t_j co-occur in W , then $s_i \leq t_j$ since W witnesses $\Phi_1 \leq \Phi_2$. The vertices corresponding to the tuples of W will be called co-occurrence vertices, and satisfy the first condition for belonging to some edge in E (that $s_i \leq t_j$). It follows from the definition of interdigitation that both $v_{1,1}$ and $v_{m,n}$ are both co-occurrence vertices. Consider a co-occurrence vertex $v_{i,j}$ not equal to $v_{m,n}$, and the lexicographically least co-occurrence vertex $v_{i',j'}$ after $v_{i,j}$ (ordering vertices by ordering the pair of subscripts). We show that i, j, i' , and j' satisfy the requirements for $\langle v_{i,j}, v_{i',j'} \rangle \in E$. If not, then either $i' > i+1$ or $j' > j+1$. If $i' > i+1$, then there can be no co-occurrence vertex $v_{i+1,j'}$, contradicting that W is piecewise total. If $j' > j+1$, then since W is piecewise total, there must be a co-occurrence vertex $v_{i'',j+1}$: but if $i'' < i$ or $i'' > i'$, this contradicts the simultaneous consistency of W , and if $i'' = i$, this contradicts the lexicographically least choice of $v_{i',j'}$. It follows that every co-occurrence vertex but $v_{m,n}$ has an edge to another co-occurrence vertex closer in Manhattan distance to $v_{m,n}$, and thus that there is a path from $v_{1,1}$ to $v_{m,n}$.

For the reverse direction assume there is a path of vertices in $SG(\Phi_1, \Phi_2)$ from $v_{1,1}$ to $v_{m,n}$ given by $v_{i_1, j_1}, v_{i_2, j_2}, \dots, v_{i_r, j_r}$ with $i_1 = j_1 = 1, i_r = m, j_r = n$. Let W be the set of state tuples corresponding to the vertices along this path. W must be simultaneously consistent with the Φ_i orderings because our directed edges are all non-decreasing in the Φ_i orderings. W must be piecewise total because no edge can cross more than one state transition in either Φ_1 or Φ_2 , by the edge set definition. So W is an interdigitation. Finally, the definition of the edge set E ensures that each tuple $\langle s_i, t_j \rangle$ in W has the property $s_i \leq t_j$, so that W is a witnessing interdigitation for $\Phi_1 \leq \Phi_2$, showing that $\Phi_1 \leq \Phi_2$, as desired. \square

Lemma 10 *Given some n , let Ψ be the conjunction of the timelines*

$$\bigcup_{i=1}^n \{(\text{PROP}_n; \text{True}_i; \text{False}_i; \text{PROP}_n), (\text{PROP}_n; \text{False}_i; \text{True}_i; \text{PROP}_n)\}.$$

We have the following facts about truth assignments to the Boolean variables p_1, \dots, p_n :

1. *For any truth assignment A , $\text{PROP}_n; s_A; \text{PROP}_n$ is semantically equivalent to some member of $\text{IS}(\Psi)$.*

2. For each $\Phi \in \text{IS}(\Psi)$ there is a truth assignment A such that $\Phi \leq \text{PROP}_n; s_A; \text{PROP}_n$.

Proof: To prove the first part of the lemma, we construct an interdigitation I of Ψ such that the corresponding member of $\text{IS}(\Psi)$ is equivalent to $\text{PROP}_n; s_A; \text{PROP}_n$. Intuitively, we construct I by ensuring that some tuple of I consists only of states of the form True_k or False_k that agree with the truth assignment—the union of all the states in this tuple, taken by $\text{IS}(\Psi)$ will equal s_A . Let $I = \{T_0, T_1, T_2, T_3, T_4\}$ be an interdigitation of Ψ with exactly five state tuples T_i . We assign the states of each timeline of Ψ to the tuples as follows:

1. For any k , such that $1 \leq k \leq n$ and $A(p_k)$ is true,
 - for the timeline $s_1; s_2; s_3; s_4 = Q; \text{True}_k; \text{False}_k; Q$, assign each state s_i to tuple T_i , and assign state s_1 to T_0 as well, and
 - for the timeline $s'_1; s'_2; s'_3; s'_4 = Q; \text{False}_k; \text{True}_k; Q$, assign each state s'_i to tuple T_{i-1} , and state s'_4 to tuple T_4 as well.
2. For any k , such that $1 \leq k \leq n$ and $A(p_k)$ is false, assign states to tuples as in item 1 while interchanging the roles of True_k and False_k .

It should be clear that I is piecewise total and simultaneously consistent with the state orderings in Ψ , and so is an interdigitation. The union of the states in each of T_0, T_1, T_3 , and T_4 is equal to PROP_n , since PROP_n is included as a state in each of those tuples. Furthermore, we see that the union of the states in T_2 is equal to s_A . Thus, the member of $\text{IS}(\Psi)$ corresponding to I is equal to $\text{PROP}_n; \text{PROP}_n; s_A; \text{PROP}_n; \text{PROP}_n$, which is semantically equivalent to $\text{PROP}_n; s_A; \text{PROP}_n$, as desired.

To prove the second part of the lemma, let Φ be any member of $\text{IS}(\Psi)$. We first argue that every state in Φ must contain either True_k or False_k for each $1 \leq k \leq n$. For any k , since Ψ contains $\text{PROP}_n; \text{True}_k; \text{False}_k; \text{PROP}_n$, every member of $\text{IS}(\Psi)$ must be subsumed by $\text{PROP}_n; \text{True}_k; \text{False}_k; \text{PROP}_n$. So, Φ is subsumed by $\text{PROP}_n; \text{True}_k; \text{False}_k; \text{PROP}_n$. But every state in $\text{PROP}_n; \text{True}_k; \text{False}_k; \text{PROP}_n$ contains either True_k or False_k , implying that so does Φ , as desired.

Next, we claim that for each $1 \leq k \leq n$, either $\Phi \leq \text{True}_k$ or $\Phi \leq \text{False}_k$ —i.e., either all states in Φ include True_k , or all states in Φ include False_k (and possibly both). To prove this claim, assume, for the sake of contradiction, that, for some k , $\Phi \not\leq \text{True}_k$ and $\Phi \not\leq \text{False}_k$. Combining this assumption with our first claim, we see there must be states s and s' in Φ such that s contains True_k but not False_k , and s' contains False_k but not True_k , respectively. Consider the interdigitation I of Ψ that corresponds to Φ (as a member of $\text{IS}(\Psi)$). We know that s and s' are each equal to the union of states in tuples T and T' , respectively, of I . T and T' must each include one state from each timeline $s_1; s_2; s_3; s_4 = \text{PROP}_n; \text{True}_k; \text{False}_k; \text{PROP}_n$ and $s'_1; s'_2; s'_3; s'_4 = \text{PROP}_n; \text{False}_k; \text{True}_k; \text{PROP}_n$. Clearly, since s does not include False_k , T includes the states s_1 and s_2 , and likewise T' includes the states s_2 and s'_1 . It follows that I is not simultaneously consistent with the state orderings in $s_1; s_2; s_3; s_4$ and $s'_1; s'_2; s'_3; s'_4$, contradicting our choice of I as an interdigitation. This shows that either $\Phi \leq \text{True}_k$ or $\Phi \leq \text{False}_k$.

Define the truth assignment A such that for all $1 \leq k \leq n$, $A(p_k)$ if and only if $\Phi \leq \text{True}_k$. Since, for each k , $\Phi \leq \text{True}_k$ or $\Phi \leq \text{False}_k$, it follows that each state of Φ is subsumed by s_A . Furthermore, since Φ begins and ends with PROP_n , it is easy to give an interdigitation of Φ and $\text{PROP}_n; s_A; \text{PROP}_n$ that witnesses $\Phi \leq \text{PROP}_n; s_A; \text{PROP}_n$. Thus, we have that $\Phi \leq \text{PROP}_n; s_A; \text{PROP}_n$. \square

Lemma 16 *Let Φ_1 and Φ_2 be as given on page 17, in the proof of Theorem 17, and let $\Psi = \bigwedge \text{IG}(\{\Phi_1, \Phi_2\})$. For any Ψ' whose timelines are a subset of those in Ψ that omits some square timeline, we have $\Psi < \Psi'$.*

Proof: Since the timelines in Ψ' are a subset of the timelines in Ψ , we know that $\Psi \leq \Psi'$. It remains to show that $\Psi' \not\leq \Psi$. We show this by constructing a timeline that is covered by Ψ' , but not by Ψ .

Let $\Phi = s_1; s_2; \dots; s_{2n-1}$ be a square timeline in Ψ that is not included in Ψ' . Recall that each s_i is a single proposition from the proposition set $P = \{p_{i,j} \mid 1 \leq i \leq n, 1 \leq j \leq n\}$, and that, for consecutive states s_i and s_{i+1} , if $s_i = p_{i,j}$, then s_{i+1} is either $p_{i+1,j}$ or $p_{i,j+1}$. Define a new timeline

$\bar{\Phi} = \bar{s}_2; \bar{s}_3; \dots; \bar{s}_{2n-2}$ with $\bar{s}_i = (P - s_i)$. We now show that $\bar{\Phi} \not\leq \Phi$ (so that $\bar{\Phi} \not\leq \Psi$), and that, for any Φ' in $\Psi - \{\Phi\}$, $\bar{\Phi} \leq \Phi'$ (so that $\bar{\Phi} \leq \Psi'$).

For the sake of contradiction, assume that $\bar{\Phi} \leq \Phi$ —then there must be a interdigitation W witnessing $\bar{\Phi} \leq \Phi$. We show by induction on i that, for $i \geq 2$, $W(s_i, \bar{s}_j)$ implies $j > i$. For the base case, when $i = 2$, we know that $\bar{s}_2 \not\leq s_2$, since $s_2 \not\leq \bar{s}_2$, and so $W(s_2, \bar{s}_2)$ is false, since W witnesses subsumption. For the inductive case, assume the claim holds for all $i' < i$, and that $W(s_i, \bar{s}_j)$. We know that $\bar{s}_i \not\leq s_i$, and thus $i \neq j$. Because W is piecewise total, we must have $W(s_{i-1}, \bar{s}_{j'})$ for some j' , and, by the induction hypothesis, we must have $j' > i-1$. Since W is simultaneously consistent with the s_k and $\bar{s}_{k'}$ state orderings, and $i-1 < i$, we have $j' \leq j$. It follows that $j > i$ as desired. Given this claim, we see that s_{2n-2} cannot co-occur in W with any state in $\bar{\Phi}$, contradicting the fact that W is piecewise total. Thus we have that $\bar{\Phi} \not\leq \Phi$.

Let $\Phi' = s'_1; \dots; s'_m$ be any timeline in $\Psi - \{\Phi\}$, we now construct an interdigitation that witnesses $\bar{\Phi} \leq \Phi'$. Note that while Φ is assumed to be square, Φ' need not be. Let j be the smallest index where $s_j \neq s'_j$ —since $s_1 = s'_1 = p_{1,1}$, and $\Phi \neq \Phi'$, we know that such a j must exist, and is in the range $2 \leq j \leq m$. We use the index j to guide our construction of an interdigitation. Let W be an interdigitation of $\bar{\Phi}$ and Φ' , with exactly the following co-occurring states (i.e., state tuples):

1. For $1 \leq i \leq j-1$, \bar{s}_{i+1} co-occurs with s'_i .
2. For $j \leq i \leq m$, \bar{s}_j co-occurs with s'_i .
3. For $j+1 \leq i \leq 2n-2$, \bar{s}_i co-occurs with s'_m .

It is easy to check that W is both piecewise total and simultaneously consistent with the state orderings in $\bar{\Phi}$ and Φ , and so is an interdigitation. We now show that W witnesses $\bar{\Phi} \leq \Phi'$ by showing that all states in $\bar{\Phi}$ are subsumed by the states they co-occur with in W . For co-occurring states \bar{s}_{i+1} and s'_i corresponding to the first item above we have that $s'_i = s_i$ —this implies that s'_i is contained in \bar{s}_{i+1} , giving that $\bar{s}_{i+1} \leq s'_i$. Now consider co-occurring states \bar{s}_j and s'_i from the second item above. Since Φ is square, choose k and l so that $s_{j-1} = p_{k,l}$, we have that s_j is either $p_{k+1,l}$ or $p_{k,l+1}$. In addition, since $s_{j-1} = s'_{j-1}$ we have that s'_j is either $p_{k+1,l}$, $p_{k,l+1}$ or $p_{k+1,l+1}$ but that $s_j \neq s'_j$. In any of these cases, we find that no state in Φ' after s'_j can equal s_j —this follows by noting that the proposition indices never decrease across the timeline Φ' . We therefore have that, for $i \geq j$, $\bar{s}_j \leq s'_i$. Finally, for co-occurring states \bar{s}_i and s'_m from item three above, we have $\bar{s}_i \leq s'_m$, since $s'_m = p_{n,n}$, which is in all states of $\bar{\Phi}$. Thus, we have shown that for all co-occurring states in W , the state from $\bar{\Phi}$ is subsumed by the co-occurring state in Φ' . Therefore, W witnesses $\bar{\Phi} \leq \Phi'$, which implies that $\bar{\Phi} \leq \Phi$. \square

7. Note that if Φ were not required to be square then it is possible for s'_{j+1} to equal s_j —i.e., they could both equal $p_{k+1,l+1}$.

Appendix C. Hand-coded and Learned Definitions Used in Our Experiments

Below we give the two sets of hand-coded definitions, HD₁ and HD₂, used in our experimental evaluation. We also give a set of learned AMA event definitions for the same seven event-types. The learned definitions correspond to the output of our k -AMA learning algorithm, given all available training examples (30 examples per event-type), with $k = 3$ and $D = NPN$. All the event definitions are written in event logic, where $\neg \diamond p$ denotes the proposition corresponding to the negation of proposition p .

$$\begin{aligned}
 \text{PICKUP}(x, y, z) &\triangleq \left(\begin{array}{l} \neg \diamond x = y \wedge \neg \diamond z = x \wedge \neg \diamond z = y \wedge \\ \text{SUPPORTED}(y) \wedge \neg \diamond \text{ATTACHED}(x, z) \wedge \\ \left[\begin{array}{l} \neg \diamond \text{ATTACHED}(x, y) \wedge \neg \diamond \text{SUPPORTS}(x, y) \wedge \\ \text{SUPPORTS}(z, y) \wedge \\ \neg \diamond \text{SUPPORTED}(x) \wedge \neg \diamond \text{ATTACHED}(y, z) \wedge \\ \neg \diamond \text{SUPPORTS}(y, x) \wedge \neg \diamond \text{SUPPORTS}(y, z) \wedge \\ \neg \diamond \text{SUPPORTS}(x, z) \wedge \neg \diamond \text{SUPPORTS}(z, x) \end{array} \right] ; \\ \left[\begin{array}{l} \text{ATTACHED}(x, y) \vee \text{ATTACHED}(y, z) ; \\ \text{ATTACHED}(x, y) \wedge \text{SUPPORTS}(x, y) \wedge \\ \neg \diamond \text{SUPPORTS}(z, y) \wedge \\ \neg \diamond \text{SUPPORTED}(x) \wedge \neg \diamond \text{ATTACHED}(y, z) \wedge \\ \neg \diamond \text{SUPPORTS}(y, x) \wedge \neg \diamond \text{SUPPORTS}(y, z) \wedge \\ \neg \diamond \text{SUPPORTS}(x, z) \wedge \neg \diamond \text{SUPPORTS}(z, x) \end{array} \right] \end{array} \right) \\
 \text{PUTDOWN}(x, y, z) &\triangleq \left(\begin{array}{l} \neg \diamond x = y \wedge \neg \diamond z = x \wedge \neg \diamond z = y \wedge \\ \text{SUPPORTED}(y) \wedge \neg \diamond \text{ATTACHED}(x, z) \wedge \\ \left[\begin{array}{l} \text{ATTACHED}(x, y) \wedge \text{SUPPORTS}(x, y) \wedge \\ \neg \diamond \text{SUPPORTS}(z, y) \wedge \\ \neg \diamond \text{SUPPORTED}(x) \wedge \neg \diamond \text{ATTACHED}(y, z) \wedge \\ \neg \diamond \text{SUPPORTS}(y, x) \wedge \neg \diamond \text{SUPPORTS}(y, z) \wedge \\ \neg \diamond \text{SUPPORTS}(x, z) \wedge \neg \diamond \text{SUPPORTS}(z, x) \end{array} \right] ; \\ \left[\begin{array}{l} \text{ATTACHED}(x, y) \vee \text{ATTACHED}(y, z) ; \\ \neg \diamond \text{ATTACHED}(x, y) \wedge \neg \diamond \text{SUPPORTS}(x, y) \wedge \\ \text{SUPPORTS}(z, y) \wedge \\ \neg \diamond \text{SUPPORTED}(x) \wedge \neg \diamond \text{ATTACHED}(y, z) \wedge \\ \neg \diamond \text{SUPPORTS}(y, x) \wedge \neg \diamond \text{SUPPORTS}(y, z) \wedge \\ \neg \diamond \text{SUPPORTS}(x, z) \wedge \neg \diamond \text{SUPPORTS}(z, x) \end{array} \right] \end{array} \right) \\
 \text{STACK}(w, x, y, z) &\triangleq \left[\begin{array}{l} \neg \diamond z = w \wedge \neg \diamond z = x \wedge \neg \diamond z = y \wedge \\ \text{PUTDOWN}(w, x, y) \wedge \text{SUPPORTS}(z, y) \wedge \\ \neg \text{ATTACHED}(z, y) \end{array} \right] \\
 \text{UNSTACK}(w, x, y, z) &\triangleq \left[\begin{array}{l} \neg \diamond z = w \wedge \neg \diamond z = x \wedge \neg \diamond z = y \wedge \\ \text{PICKUP}(w, x, y) \wedge \text{SUPPORTS}(z, y) \wedge \neg \text{ATTACHED}(z, y) \end{array} \right] \\
 \text{MOVE}(w, x, y, z) &\triangleq \neg \diamond y = z \wedge [\text{PICKUP}(w, x, y); \text{PUTDOWN}(w, x, z)] \\
 \text{ASSEMBLE}(w, x, y, z) &\triangleq \text{PUTDOWN}(w, y, z) \wedge_{\{<\}} \text{STACK}(w, x, y, z) \\
 \text{DISASSEMBLE}(w, x, y, z) &\triangleq \text{UNSTACK}(w, x, y, z) \wedge_{\{<\}} \text{PICKUP}(x, y, z)
 \end{aligned}$$

Figure 10: The HD₁ event logic definitions for all seven event-types.

$$\begin{array}{l}
 \text{PICKUP}(x, y, z) \triangleq \left(\begin{array}{l}
 \neg\Diamond x = y \wedge \neg\Diamond z = x \wedge \neg\Diamond z = y \wedge \\
 \text{SUPPORTED}(y) \wedge \neg\Diamond \text{ATTACHED}(x, z) \wedge \\
 \left[\begin{array}{l}
 \neg\Diamond \text{ATTACHED}(x, y) \wedge \neg\Diamond \text{SUPPORTS}(x, y) \wedge \\
 \text{SUPPORTS}(z, y) \wedge \text{CONTACTS}(z, y) \wedge \\
 \neg\Diamond \text{SUPPORTED}(x) \wedge \neg\Diamond \text{ATTACHED}(y, z) \wedge \\
 \neg\Diamond \text{SUPPORTS}(y, x) \wedge \neg\Diamond \text{SUPPORTS}(y, z) \wedge \\
 \neg\Diamond \text{SUPPORTS}(x, z) \wedge \neg\Diamond \text{SUPPORTS}(z, x)
 \end{array} \right] \wedge_{\{<,m\}} \\
 \left[\begin{array}{l}
 \text{ATTACHED}(x, y) \wedge \text{SUPPORTS}(x, y) \wedge \\
 \neg\Diamond \text{SUPPORTS}(z, y) \wedge \\
 \neg\Diamond \text{SUPPORTED}(x) \wedge \neg\Diamond \text{ATTACHED}(y, z) \wedge \\
 \neg\Diamond \text{SUPPORTS}(y, x) \wedge \neg\Diamond \text{SUPPORTS}(y, z) \wedge \\
 \neg\Diamond \text{SUPPORTS}(x, z) \wedge \neg\Diamond \text{SUPPORTS}(z, x)
 \end{array} \right]
 \end{array} \right) \\
 \\
 \text{PUTDOWN}(x, y, z) \triangleq \left(\begin{array}{l}
 \neg\Diamond x = y \wedge \neg\Diamond z = x \wedge \neg\Diamond z = y \wedge \\
 \text{SUPPORTED}(y) \wedge \neg\Diamond \text{ATTACHED}(x, z) \wedge \\
 \left[\begin{array}{l}
 \text{ATTACHED}(x, y) \wedge \text{SUPPORTS}(x, y) \wedge \\
 \neg\Diamond \text{SUPPORTS}(z, y) \wedge \\
 \neg\Diamond \text{SUPPORTED}(x) \wedge \neg\Diamond \text{ATTACHED}(y, z) \wedge \\
 \neg\Diamond \text{SUPPORTS}(y, x) \wedge \neg\Diamond \text{SUPPORTS}(y, z) \wedge \\
 \neg\Diamond \text{SUPPORTS}(x, z) \wedge \neg\Diamond \text{SUPPORTS}(z, x)
 \end{array} \right] \wedge_{\{<,m\}} \\
 \left[\begin{array}{l}
 \neg\Diamond \text{ATTACHED}(x, y) \wedge \neg\Diamond \text{SUPPORTS}(x, y) \wedge \\
 \text{SUPPORTS}(z, y) \wedge \text{CONTACTS}(z, y) \wedge \\
 \neg\Diamond \text{SUPPORTED}(x) \wedge \neg\Diamond \text{ATTACHED}(y, z) \wedge \\
 \neg\Diamond \text{SUPPORTS}(y, x) \wedge \neg\Diamond \text{SUPPORTS}(y, z) \wedge \\
 \neg\Diamond \text{SUPPORTS}(x, z) \wedge \neg\Diamond \text{SUPPORTS}(z, x)
 \end{array} \right]
 \end{array} \right)
 \end{array}$$

 Figure 11: Part I of the HD₂ event logic definitions.

$$\begin{aligned}
 \text{STACK}(w, x, y, z) &\triangleq \left(\begin{array}{l} \neg\Diamond w = x \wedge \neg\Diamond y = w \wedge \neg\Diamond y = x \wedge \\ \neg\Diamond z = w \wedge \neg\Diamond z = x \wedge \neg\Diamond z = y \wedge \\ \text{SUPPORTED}(x) \wedge \neg\Diamond \text{ATTACHED}(w, y) \wedge \\ \left[\begin{array}{l} \text{ATTACHED}(w, x) \wedge \text{SUPPORTS}(w, x) \wedge \\ \neg\Diamond \text{SUPPORTS}(y, x) \wedge \\ \text{SUPPORTS}(z, y) \wedge \text{CONTACTS}(z, y) \wedge \\ \neg\Diamond \text{ATTACHED}(z, y) \wedge \\ \neg\Diamond \text{SUPPORTED}(w) \wedge \neg\Diamond \text{ATTACHED}(x, y) \wedge \\ \neg\Diamond \text{SUPPORTS}(x, w) \wedge \neg\Diamond \text{SUPPORTS}(x, y) \wedge \\ \neg\Diamond \text{SUPPORTS}(w, y) \wedge \neg\Diamond \text{SUPPORTS}(y, w) \end{array} \right] \wedge_{\{<,m\}} \\ \left[\begin{array}{l} \neg\Diamond \text{ATTACHED}(w, x) \wedge \neg\Diamond \text{SUPPORTS}(w, x) \wedge \\ \text{SUPPORTS}(y, x) \wedge \text{CONTACTS}(y, x) \wedge \\ \text{SUPPORTS}(z, y) \wedge \text{CONTACTS}(z, y) \wedge \\ \neg\Diamond \text{ATTACHED}(z, y) \wedge \\ \neg\Diamond \text{SUPPORTED}(w) \wedge \neg\Diamond \text{ATTACHED}(x, y) \wedge \\ \neg\Diamond \text{SUPPORTS}(x, w) \wedge \neg\Diamond \text{SUPPORTS}(x, y) \wedge \\ \neg\Diamond \text{SUPPORTS}(w, y) \wedge \neg\Diamond \text{SUPPORTS}(y, w) \end{array} \right] \end{array} \right) \\
 \\
 \text{UNSTACK}(w, x, y, z) &\triangleq \left(\begin{array}{l} \neg\Diamond w = x \wedge \neg\Diamond y = w \wedge \neg\Diamond y = x \wedge \\ \neg\Diamond z = w \wedge \neg\Diamond z = x \wedge \neg\Diamond z = y \wedge \\ \text{SUPPORTED}(x) \wedge \neg\Diamond \text{ATTACHED}(w, y) \wedge \\ \left[\begin{array}{l} \neg\Diamond \text{ATTACHED}(w, x) \wedge \neg\Diamond \text{SUPPORTS}(w, x) \wedge \\ \text{SUPPORTS}(y, x) \wedge \text{CONTACTS}(y, x) \wedge \\ \text{SUPPORTS}(z, y) \wedge \text{CONTACTS}(z, y) \wedge \\ \neg\Diamond \text{ATTACHED}(z, y) \wedge \\ \neg\Diamond \text{SUPPORTED}(w) \wedge \neg\Diamond \text{ATTACHED}(x, y) \wedge \\ \neg\Diamond \text{SUPPORTS}(x, w) \wedge \neg\Diamond \text{SUPPORTS}(x, y) \wedge \\ \neg\Diamond \text{SUPPORTS}(w, y) \wedge \neg\Diamond \text{SUPPORTS}(y, w) \end{array} \right] \wedge_{\{<,m\}} \\ \left[\begin{array}{l} \text{ATTACHED}(w, x) \wedge \text{SUPPORTS}(w, x) \wedge \\ \neg\Diamond \text{SUPPORTS}(y, x) \wedge \\ \text{SUPPORTS}(z, y) \wedge \text{CONTACTS}(z, y) \wedge \\ \neg\Diamond \text{ATTACHED}(z, y) \wedge \\ \neg\Diamond \text{SUPPORTED}(w) \wedge \neg\Diamond \text{ATTACHED}(x, y) \wedge \\ \neg\Diamond \text{SUPPORTS}(x, w) \wedge \neg\Diamond \text{SUPPORTS}(x, y) \wedge \\ \neg\Diamond \text{SUPPORTS}(w, y) \wedge \neg\Diamond \text{SUPPORTS}(y, w) \end{array} \right] \end{array} \right) \\
 \\
 \text{MOVE}(w, x, y, z) &\triangleq \neg\Diamond y = z \wedge [\text{PICKUP}(w, x, y); \text{PUTDOWN}(w, x, z)] \\
 \text{ASSEMBLE}(w, x, y, z) &\triangleq \text{PUTDOWN}(w, y, z) \wedge_{\{<\}} \text{STACK}(w, x, y, z) \\
 \text{DISASSEMBLE}(w, x, y, z) &\triangleq \text{UNSTACK}(w, x, y, z) \wedge_{\{<\}} \text{PICKUP}(x, y, z)
 \end{aligned}$$

 Figure 12: Part II of the HD₂ event logic definitions.

$$\begin{aligned}
 \text{PICKUP}(x, y, z) &\triangleq \left(\left\{ \begin{array}{l} \text{SUPPORTED}(y); \\ \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \text{ATTACHED}(x, y) \wedge \\ \text{ATTACHED}(y, z) \end{array} \right]; \end{array} \right\} \wedge \right. \\
 &\left. \left\{ \begin{array}{l} [\text{SUPPORTED}(y) \wedge \text{ATTACHED}(x, y)] \\ [\text{SUPPORTED}(y) \wedge \text{CONTACTS}(y, z)]; \\ [\text{SUPPORTED}(y) \wedge \text{ATTACHED}(y, z)]; \\ [\text{SUPPORTED}(y) \wedge \text{ATTACHED}(x, y)] \end{array} \right\} \wedge \right. \\
 &\left. \left\{ \begin{array}{l} \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(z, y) \wedge \\ \text{CONTACTS}(y, z) \wedge \neg \diamond \text{SUPPORTS}(x, y) \wedge \\ \neg \diamond \text{ATTACHED}(x, y) \wedge \neg \diamond \text{ATTACHED}(y, z) \end{array} \right]; \\ [\text{SUPPORTED}(y) \wedge \text{SUPPORTS}(z, y)]; \\ [\text{SUPPORTED}(y) \wedge \text{ATTACHED}(x, y)] \end{array} \right\} \wedge \right. \\
 &\left. \left\{ \begin{array}{l} \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(z, y) \wedge \\ \text{CONTACTS}(y, z) \wedge \neg \diamond \text{SUPPORTS}(x, y) \wedge \\ \neg \diamond \text{ATTACHED}(x, y) \wedge \neg \diamond \text{ATTACHED}(y, z) \end{array} \right]; \\ \text{SUPPORTED}(y); \\ \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(x, y) \wedge \\ \text{ATTACHED}(x, y) \wedge \neg \diamond \text{SUPPORTS}(z, y) \wedge \\ \neg \diamond \text{CONTACTS}(y, z) \wedge \neg \diamond \text{ATTACHED}(y, z) \end{array} \right]; \\ [\text{SUPPORTED}(y) \wedge \text{SUPPORTS}(z, y)]; \\ [\text{SUPPORTED}(y) \wedge \text{ATTACHED}(x, y)]; \\ \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(x, y) \wedge \\ \text{ATTACHED}(x, y) \wedge \neg \diamond \text{SUPPORTS}(z, y) \wedge \\ \neg \diamond \text{CONTACTS}(y, z) \wedge \neg \diamond \text{ATTACHED}(y, z) \end{array} \right] \end{array} \right\} \wedge \right. \\
 &\left. \left\{ \begin{array}{l} [\text{SUPPORTED}(y) \wedge \text{ATTACHED}(x, y)]; \\ [\text{SUPPORTED}(y) \wedge \text{ATTACHED}(x, y) \wedge \text{ATTACHED}(y, z)]; \\ \text{SUPPORTED}(y) \end{array} \right\} \wedge \right. \\
 &\left. \left\{ \begin{array}{l} \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \text{ATTACHED}(x, y) \wedge \text{ATTACHED}(x, y) \wedge \\ \neg \diamond \text{SUPPORTS}(z, y) \wedge \neg \diamond \text{CONTACTS}(y, z) \wedge \neg \diamond \text{ATTACHED}(y, z) \end{array} \right]; \\ \text{SUPPORTED}(y); \\ \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(x, y) \wedge \text{CONTACTS}(z, y) \wedge \\ \neg \diamond \text{SUPPORTS}(x, y) \wedge \neg \diamond \text{ATTACHED}(x, y) \end{array} \right] \end{array} \right\} \right) \\
 \\
 \text{PUTDOWN}(x, y, z) &\triangleq \left(\left\{ \begin{array}{l} \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \text{ATTACHED}(x, y) \end{array} \right]; \\ \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \text{ATTACHED}(x, y) \wedge \text{ATTACHED}(y, z) \end{array} \right]; \\ \text{SUPPORTED}(y) \end{array} \right\} \wedge \right. \\
 &\left. \left\{ \begin{array}{l} \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \text{ATTACHED}(x, y) \wedge \text{ATTACHED}(x, y) \wedge \\ \neg \diamond \text{SUPPORTS}(z, y) \wedge \neg \diamond \text{CONTACTS}(y, z) \wedge \neg \diamond \text{ATTACHED}(y, z) \end{array} \right]; \\ \text{SUPPORTED}(y); \\ \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(x, y) \wedge \text{CONTACTS}(z, y) \wedge \\ \neg \diamond \text{SUPPORTS}(x, y) \wedge \neg \diamond \text{ATTACHED}(x, y) \end{array} \right] \end{array} \right\} \right)
 \end{aligned}$$

Figure 13: The learned 3-AMA definitions for PICKUP(x,y,z) and PUTDOWN(x,y,z).

$$\left(\left\{ \begin{array}{l} [\text{SUPPORTED}(x) \wedge \text{SUPPORTS}(y, x)]; \\ [\text{SUPPORTED}(x) \wedge \text{ATTACHED}(w, x)]; \\ \text{SUPPORTED}(x) \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{SUPPORTED}(x); \\ [\text{SUPPORTED}(x) \wedge \text{ATTACHED}(w, x) \wedge \text{ATTACHED}(x, z)]; \\ \text{SUPPORTED}(x) \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} [\text{SUPPORTED}(x)]; \\ [\text{SUPPORTED}(x) \wedge \text{ATTACHED}(x, z)]; \\ [\text{SUPPORTED}(x) \wedge \text{CONTACTS}(x, z)] \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{SUPPORTED}(x); \\ [\text{SUPPORTED}(x) \wedge \text{ATTACHED}(w, x) \wedge \text{SUPPORTS}(w, x)]; \\ \text{SUPPORTED}(x) \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{SUPPORTED}(x); \\ [\text{SUPPORTED}(x) \wedge \text{ATTACHED}(w, x) \wedge \text{ATTACHED}(y, x)]; \\ \text{SUPPORTED}(x) \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} [\text{SUPPORTED}(x) \wedge \text{CONTACTS}(y, x)]; \\ [\text{SUPPORTED}(x) \wedge \text{ATTACHED}(y, x)]; \\ \text{SUPPORTED}(x) \end{array} \right\} \wedge \right. \\
 \left. \left[\begin{array}{l} \text{SUPPORTED}(x) \wedge \text{SUPPORTS}(y, x) \wedge \text{CONTACTS}(y, x) \wedge \\ \neg \diamond \text{SUPPORTS}(w, x) \wedge \neg \diamond \text{SUPPORTS}(z, x) \wedge \neg \diamond \text{CONTACTS}(x, z) \wedge \\ \neg \diamond \text{ATTACHED}(w, x) \wedge \neg \diamond \text{ATTACHED}(y, x) \wedge \neg \diamond \text{ATTACHED}(x, z) \end{array} \right]; \right. \\
 \left. \left[\begin{array}{l} \text{SUPPORTED}(x); \\ \text{SUPPORTED}(x) \wedge \text{SUPPORTS}(z, x) \wedge \text{CONTACTS}(x, z) \wedge \\ \neg \diamond \text{SUPPORTS}(w, x) \wedge \neg \diamond \text{SUPPORTS}(y, x) \wedge \neg \diamond \text{CONTACTS}(y, x) \wedge \\ \neg \diamond \text{ATTACHED}(w, x) \wedge \neg \diamond \text{ATTACHED}(y, x) \wedge \neg \diamond \text{ATTACHED}(x, z) \end{array} \right] \right)$$

Figure 16: The learned 3-AMA definition for MOVE(w,x,y,z).

$$\left(\left\{ \left[\begin{array}{l} \neg\Diamond\text{SUPPORTED}(x) \wedge \neg\Diamond\text{SUPPORTS}(z, y) \wedge \neg\Diamond\text{SUPPORTS}(y, x) \wedge \neg\Diamond\text{CONTACTS}(x, Y) \wedge \\ \neg\Diamond\text{CONTACTS}(z, y) \wedge \neg\Diamond\text{ATTACHED}(w, x) \wedge \neg\Diamond\text{ATTACHED}(z, y) \end{array} \right]; \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{ATTACHED}(w, y); \\ \text{SUPPORTED}(y) \end{array} \right\} \wedge \right. \\
 \left. \left[\begin{array}{l} \neg\Diamond\text{SUPPORTED}(x) \wedge \neg\Diamond\text{SUPPORTS}(z, y) \wedge \neg\Diamond\text{SUPPORTS}(y, x) \wedge \neg\Diamond\text{CONTACTS}(x, Y) \wedge \\ \neg\Diamond\text{CONTACTS}(z, y) \wedge \neg\Diamond\text{ATTACHED}(w, x) \wedge \neg\Diamond\text{ATTACHED}(z, y) \end{array} \right]; \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{true}; \\ \left[\begin{array}{l} \text{SUPPORTED}(x) \wedge \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(z, y) \wedge \text{SUPPORTS}(y, x) \wedge \\ \text{CONTACTS}(x, y) \wedge \text{CONTACTS}(z, y) \wedge \neg\Diamond\text{ATTACHED}(w, y) \end{array} \right] \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{true}; \\ [\text{SUPPORTED}(y) \wedge \neg\Diamond\text{ATTACHED}(w, x) \wedge \Diamond\text{ATTACHED}(z, y)]; \\ \text{SUPPORTED}(y) \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{true}; \\ [\text{SUPPORTED}(y) \wedge \text{ATTACHED}(z, y)]; \\ [\text{SUPPORTED}(y) \wedge \text{CONTACTS}(z, y)] \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{true}; \\ [\text{SUPPORTED}(y) \wedge \text{SUPPORTS}(z, y) \text{CONTACTS}(z, y) \wedge \text{ATTACHED}(w, x)]; \\ \text{SUPPORTED}(y) \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{true}; \\ [\text{SUPPORTED}(y) \wedge \text{ATTACHED}(w, y) \text{ATTACHED}(z, y)]; \\ \text{SUPPORTED}(y) \end{array} \right\} \right)$$

Figure 17: The learned 3-AMA definition for ASSEMBLE(w,x,y,z).

$$\left(\left\{ \begin{array}{l} \text{SUPPORTED}(x) \wedge \text{SUPPORTED}(y); \\ \left[\begin{array}{l} \text{SUPPORTED}(x) \wedge \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(w, x) \wedge \\ \text{SUPPORTS}(z, y) \wedge \text{CONTACTS}(z, y) \wedge \text{ATTACHED}(x, w) \end{array} \right]; \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{SUPPORTED}(y) \\ \left[\begin{array}{l} \text{SUPPORTED}(x) \wedge \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(z, y) \wedge \\ \text{SUPPORTS}(y, x) \wedge \text{CONTACTS}(x, y) \wedge \text{CONTACTS}(z, y) \end{array} \right]; \\ \text{SUPPORTED}(x) \wedge \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(y, x) \wedge \text{ATTACHED}(x, y); \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{SUPPORTED}(y) \\ \left[\begin{array}{l} \text{SUPPORTED}(x) \wedge \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(y, x) \wedge \text{CONTACTS}(z, y); \\ \left[\begin{array}{l} \text{SUPPORTED}(x) \wedge \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(x, y) \wedge \\ \text{SUPPORTS}(y, z) \wedge \text{ATTACHED}(x, y) \wedge \text{ATTACHED}(z, y) \end{array} \right]; \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{SUPPORTED}(y) \\ \left[\begin{array}{l} \text{SUPPORTED}(x) \wedge \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(y, x); \\ \left[\begin{array}{l} \text{SUPPORTED}(x) \wedge \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(x, y) \wedge \\ \text{SUPPORTS}(y, z) \wedge \text{ATTACHED}(x, y) \wedge \text{ATTACHED}(z, y) \wedge \text{ATTACHED}(x, w) \end{array} \right]; \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{SUPPORTED}(y); \\ \left[\begin{array}{l} \text{SUPPORTED}(x) \wedge \text{SUPPORTED}(y) \wedge \text{SUPPORTS}(y, x) \wedge \text{SUPPORTS}(z, y) \wedge \\ \text{CONTACTS}(x, y) \wedge \text{CONTACTS}(z, y) \wedge \neg \diamond \text{SUPPORTS}(w, x) \wedge \\ \neg \diamond \text{SUPPORTS}(w, y) \wedge \neg \diamond \text{SUPPORTS}(x, y) \wedge \neg \diamond \text{ATTACHED}(x, w) \wedge \\ \neg \diamond \text{ATTACHED}(w, y) \wedge \neg \diamond \text{ATTACHED}(x, y) \wedge \neg \diamond \text{ATTACHED}(z, y) \end{array} \right]; \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{SUPPORTED}(y); \\ \left[\begin{array}{l} \text{SUPPORTED}(y) \wedge \neg \diamond \text{SUPPORTED}(x) \wedge \neg \diamond \text{SUPPORTS}(w, x) \wedge \\ \neg \diamond \text{SUPPORTS}(z, y) \wedge \neg \diamond \text{SUPPORTS}(y, x) \wedge \neg \diamond \text{CONTACTS}(x, y) \wedge \\ \neg \diamond \text{CONTACTS}(z, y) \wedge \neg \diamond \text{ATTACHED}(x, w) \wedge \neg \diamond \text{ATTACHED}(z, y) \end{array} \right]; \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{SUPPORTED}(y); \\ \left[\text{SUPPORTED}(y) \wedge \text{ATTACHED}(w, y) \wedge \text{ATTACHED}(z, y) \right]; \end{array} \right\} \wedge \right. \\
 \left. \left\{ \begin{array}{l} \text{SUPPORTED}(y); \\ \left[\text{SUPPORTED}(y); \right. \\ \left. \left[\text{SUPPORTED}(y) \wedge \text{SUPPORTS}(w, y) \wedge \text{ATTACHED}(w, y) \right]; \right. \\ \left. \text{SUPPORTED}(y) \right\} \right)
 \end{array}$$

Figure 18: The learned 3-AMA definition for DISASSEMBLE(w,x,y,z).