

A Probabilistic Logic Programming Event Calculus

Anastasios Skarlatidis^{1,2}, Alexander Artikis¹, Jason Filippou^{1,3} and Georgios Paliouras¹

¹*Institute of Informatics and Telecommunications, NCSR Demokritos, Athens, Greece*

²*Department of Digital Systems, University of Piraeus, Greece*

³*University of Maryland, USA*

(*e-mail: {anskar1, a.artikis, jfilip, paliourg}@iit.demokritos.gr*)

submitted 06 April 2012; revised 10 November 2012; accepted 07 March 2013

Note: The paper has been accepted for publication in the *Theory and Practice of Logic Programming (TPLP) Journal*.

Abstract

We present a system for recognising human activity given a symbolic representation of video content. The input of our system is a set of time-stamped short-term activities (STA) detected on video frames. The output is a set of recognised long-term activities (LTA), which are pre-defined temporal combinations of STA. The constraints on the STA that, if satisfied, lead to the recognition of a LTA, have been expressed using a dialect of the Event Calculus. In order to handle the uncertainty that naturally occurs in human activity recognition, we adapted this dialect to a state-of-the-art probabilistic logic programming framework. We present a detailed evaluation and comparison of the crisp and probabilistic approaches through experimentation on a benchmark dataset of human surveillance videos.

KEYWORDS: Event Recognition, Event Pattern Matching, Event Calculus, ProbLog

1 Introduction

Systems for event recognition — ‘event pattern matching’, in the terminology of (Luckham 2002) — accept as input streams of sensor data in order to identify composite events of interest, that is, collections of events that satisfy some pattern. Consider, for example, the recognition of attacks on nodes of a computer network given the TCP/IP messages, the recognition of suspicious trader behaviour given the transactions in a financial market, the recognition of whale songs given a stream of whale sounds, and the recognition of human activities given multimedia content from surveillance cameras.

A common approach to event recognition separates low-level from high-level recognition. In the case of human activity recognition, the output of the former type of recognition is a set of activities taking place in a short period of time: ‘short-term activities’ (STA). The output of the latter type of recognition is a set of ‘long-term activities’ (LTA), which are temporal combinations of STA. We focus on high-level recognition.

We define a set of LTA of interest, such as ‘fighting’ and ‘meeting’, as temporal combinations of STA such as ‘walking’, ‘running’, and ‘inactive’ (standing still) using a logic programming (Prolog) implementation of the Event Calculus (EC) (Kowalski and Sergot 1986). We employ EC to express the temporal constraints on a set of STA that, if satisfied, lead to the recognition of a LTA.

In earlier work (Artikis, Sergot and Paliouras (2010)) we identified various types of uncertainty that exist in activity recognition, such as erroneous STA detection. To address this issue, we extend our work by presenting Prob-EC, an EC dialect suitable for probabilistic activity recognition. Prob-EC operates on the state-of-the-art probabilistic logic programming framework ProbLog (Kimmig et al. 2011). Prob-EC, therefore, may operate in settings where STA occurrences are assigned a confidence value by the underlying low-level tracking system (such as, for example, a probabilistic classifier). We present extensive experimental evaluation of Prob-EC on a benchmark activity recognition dataset. The evaluation demonstrates the conditions in which Prob-EC outperforms our previous EC dialect — Crisp-EC. Prob-EC is the first EC dialect able to deal with uncertainty in the input STA. Moreover, this is the first approach that thoroughly evaluates EC in a probabilistic framework. The full code of Prob-EC, along with the dataset on which experimentation is performed, is available upon request.

The remainder of the paper is organised as follows. In the following section we set our work in context. In Section 3 we present Crisp-EC while in Sections 4 and 5 we describe, respectively, the dataset on which we perform activity recognition and the corresponding knowledge base of LTA definitions. Section 6 briefly introduces ProbLog while Section 7 describes Prob-EC. Our experimental results, including the comparison between Prob-EC and Crisp-EC, are presented in Section 8. Finally, in Section 9 we summarise our observations and outline directions for further research.

2 Related Work

Numerous recognition systems have been proposed in the literature (Cugola and Margara 2012). In this section we focus on long-term activity (LTA) recognition systems that, similar to our approach, exhibit a formal, declarative semantics.

A fair amount of recognition systems is logic-based. Notable approaches include the Chronicle Recognition System (Dousson and Maigat 2007) and the hierarchical event representation of (Hakeem and Shah 2007). A recent review of logic-based recognition systems may be found in (Artikis, Skarlatidis et al. (2012)). These systems are common in that they employ logic-based methods for representation and inference, but are unable to handle noise.

Shet et al. (2005; 2007) have presented a logic programming approach to activity recognition which touches upon the issue of data coming from noisy sensors. In that work, LTA concerning theft, entry violation, unattended packages, and so on, have been defined. Within their activity recognition system, Shet and colleagues have incorporated mechanisms for reasoning over rules and facts that have an uncertainty value attached. Uncertainty in rules corresponds to a measure of rule reliability. On the other hand, uncertainty in facts represents the detection probabilities of the

short-term activities (STA). In the VidMAP system (Shet et al. 2005), a mid-level module which generates Prolog facts automatically filters out data that a low-level image processing module has misclassified (such as a tree mistaken for a human). Shet and colleagues have noted of the filtering carried out by this module that ‘...it does so by observing whether or not the object has been persistently tracked’ (Shet et al. 2005, p. 2). In (Shet et al. 2007), the authors use an algebraic data structure known as a bilattice (Ginsberg 1990) to detect human entities based on uncertain output of part-based detectors, such as head or leg detectors. The bilattice structure associates every STA or LTA with two uncertainty values, one encoding available information and the other encoding confidence. The more confident information is provided, the more probable the respective LTA becomes.

A ProbLog-based method for robotic action recognition is proposed in (Moldovan et al. 2012). The method employs a relational extension of the *affordance models* (Gibson 1979) in order to represent multi-object interactions in a scene. Affordances can model the relations between objects, actions (that is, pre-programmed robotic arm movements) and effects of actions. In contrast to a standard propositional Bayesian Network implementation of an affordance model, the method can scale to multiple object interactions in a scene without the need of retraining. However, the proposed method does not include temporal representation and reasoning.

Probabilistic graphical models have been applied on a variety of activity recognition applications where uncertainty exists. Activity recognition requires processing streams of timestamped STA and, therefore, numerous activity recognition methods are based on sequential variants of probabilistic graphical models, such as Hidden Markov Models (Rabiner and Juang 1986), Dynamic Bayesian Networks (Murphy 2002) and Conditional Random Fields (Lafferty et al. 2001). Compared to logic-based methods, graphical models can naturally handle uncertainty but their propositional structure provides limited representation capabilities. To model LTA that involve a large number of relations among STA, such as interactions between multiple persons and/or objects, the structure of the model may become prohibitively large and complex. To overcome such limitations, these models have been extended in order to support more complex relations. Examples of such extensions include representing interactions involving multiple domain objects (Brand et al. 1997; Gong and Xiang 2003; Wu et al. 2007; Vail et al. 2007), capturing long-term dependencies between states (Hongeng and Nevatia 2003), as well as describing a hierarchical composition of activities (Natarajan and Nevatia 2007; Liao et al. 2007). However, the lack of a formal representation language makes the definition of complex LTA complicated and the integration of domain background knowledge very hard.

Markov Logic Networks (MLN) (Richardson and Domingos 2006) have also been used for representing uncertainty in activity recognition. MLN employ first-order logic representation, where each formula may be associated with a weight, indicating the confidence we have on the formula. The knowledge base of weighted formulas is translated into a Markov network where probabilistic inference is performed. The approach of (Biswas et al. 2007), for example, uses MLN to recognise LTA given the STA that have been observed by low-level classifiers. A more expressive approach

that can represent persistent and concurrent LTA, as well as their starting and ending time-points, is proposed in (Helaoui et al. 2011). The method in (Sadilek and Kautz 2012) employs hybrid-MLN (Wang and Domingos 2008) in order to recognise successful and failed interactions between multiple humans using noisy location data. Similar to pure MLN-based methods, the knowledge base is composed by LTA definitions. Furthermore, hybrid formulas aiming to remove the noise from the location data are also included. Hybrid formulas are defined as normal formulas, but their weights are also associated with a real-valued function, such as the distance of two persons. As a result, the confidence of the formula is defined by both its weight and function. Although these methods incorporate first-order logic representation, the presented LTA definitions have a limited temporal representation.

A method that uses interval-based temporal relations is proposed in (Morariu and Davis 2011). The aim of the method is to determine the most consistent sequence of LTA based on the observations of low-level classifiers. Similar to (Tran and Davis 2008; Kembhavi et al. 2010), the method uses MLN to express LTA. In contrast to (Tran and Davis 2008; Kembhavi et al. 2010), it employs temporal relations based on Allen’s Interval Algebra (IA) (Allen 1983). In order to avoid the combinatorial explosion of possible intervals that IA may produce, a bottom-up process eliminates the unlikely LTA hypotheses. In (Brendel et al. 2011; Selman et al. 2011) a probabilistic extension of Event Logic (Siskind 2001) is proposed in order to perform interval-based activity recognition. Similar to MLN, the method defines a probabilistic model from a set of domain-specific weighted LTA. However, the Event Logic representation avoids the enumeration of all possible interval relations.

The main difference of our approach with respect to the aforementioned lines of work concerns the fact that we use the Event Calculus (EC) for temporal representation. EC has built-in rules for complex temporal representation, including the formalisation of inertia, which help considerably the system designer in developing activity definitions. With the use of EC one may develop intuitive, succinct activity definitions, facilitating the interaction between activity definition developer and domain expert, and allowing for code maintenance. Furthermore, being logic programming-based, EC has direct routes to the ProbLog probabilistic logic programming framework. To the best of our knowledge, the probabilistic EC dialect presented in this paper, Prob-EC, is the first EC dialect able to deal with uncertainty in STA detection.

A MLN-based approach that is complementary to our work is that of (Skarlatidis et al. 2011), which introduces a probabilistic EC dialect based on MLN. This dialect and Prob-EC tackle the problem of probabilistic inference from different viewpoints. Prob-EC handles noise in the input stream, represented as detection probabilities of the STA. The MLN-based EC dialect, on the other hand, emphasises uncertainty in activity definitions in the form of rule weights.

ProbLog and MLN are closely related. A notable difference between them is that MLN, as an extension of first-order logic, are not bound by the closed-world assumption. There exists a body of work that investigates the connection between the two frameworks. (Bruynooghe et al. 2010), for example, developed an extension

Predicate	Meaning
<code>happensAt(E, T)</code>	Event E is occurring at time T
<code>initially($F = V$)</code>	The value of fluent F is V at time 0
<code>holdsAt($F = V, T$)</code>	The value of fluent F is V at time T
<code>initiatedAt($F = V, T$)</code>	At time T a period of time for which $F = V$ is initiated
<code>terminatedAt($F = V, T$)</code>	At time T a period of time for which $F = V$ is terminated

Table 1: Main Predicates of Crisp-EC.

of ProbLog which is able to handle first-order formulas with weighted constraints. Fierens et. al (2011) converted probabilistic logic programs to ground MLN and then used state-of-the-art MLN inference algorithms to perform inference on the transformed programs.

3 The Event Calculus

Our LTA recognition system is based on a logic programming (Prolog) implementation of an EC dialect. EC, introduced by Kowalski and Sergot (1986), is a many-sorted, first-order predicate calculus for representing and reasoning about events and their effects. For the dialect presented here — Crisp-EC — the time model is linear and includes integers. Where F is a *fluent* — a property that is allowed to have different values at different points in time — the term $F = V$ denotes that fluent F has value V . Boolean fluents are a special case in which the possible values are `true` and `false`. Informally, $F = V$ holds at a particular time-point if $F = V$ has been *initiated* by an event at some earlier time-point, and not *terminated* by another event in the meantime — law of inertia.

We represent STA as events and LTA as fluents. In this way, we can express the conditions in which the occurrence of a STA initiates or terminates a LTA.

An *event description* in Crisp-EC includes rules that define, among other things, the event occurrences (with the use of the `happensAt` predicate), the effects of events (with the use of the `initiatedAt` and `terminatedAt` predicates), and the values of the fluents (with the use of the `initially` and `holdsAt` predicates). Table 1 summarises the main predicates of Crisp-EC. Variables, starting with an upper-case letter, are assumed to be universally quantified unless otherwise indicated. Predicates, function symbols and constants start with a lower-case letter.

The domain-independent rules for `holdsAt` can be written in the following form:

$$\begin{aligned}
 \text{holdsAt}(F = V, T) \leftarrow & \\
 & \text{initially}(F = V), \\
 & \text{not broken}(F = V, 0, T)
 \end{aligned} \tag{1}$$

$$\begin{aligned}
\text{holdsAt}(F = V, T) \leftarrow & \\
& \text{initiatedAt}(F = V, T_s), \\
& T_s < T, \\
& \text{not broken}(F = V, T_s, T)
\end{aligned} \tag{2}$$

$$\begin{aligned}
\text{broken}(F = V, T_s, T) \leftarrow & \\
& \text{terminatedAt}(F = V, T_f), \\
& T_s < T_f < T
\end{aligned} \tag{3}$$

$$\begin{aligned}
\text{broken}(F = V_1, T_s, T) \leftarrow & \\
& \text{initiatedAt}(F = V_2, T_f), \\
& V_1 \neq V_2, \\
& T_s < T_f < T
\end{aligned} \tag{4}$$

not represents ‘negation by failure’, which provides a form of default persistence — inertia — of fluents. According to rule (1), $F = V$ holds at time-point T if $F = V$ held initially and has not been broken since. According to rule (2), $F = V$ holds at time-point T if the fluent F has been initiated to value V at an earlier time T_s , and has not been broken since. According to rule (3), a period of time for which $F = V$ holds is broken at T_f if $F = V$ is terminated at T_f . Rule (4) dictates that if $F = V_2$ is initiated at T_f then effectively $F = V_1$ is terminated at T_f , for all other possible values V_1 of F . Rule (4) therefore ensures that a fluent cannot have more than one value at any time. We do not insist that a fluent must have a value at every time-point. In Crisp-EC there is a difference between initiating a Boolean fluent $F = \text{false}$ and terminating $F = \text{true}$: the first implies, but is not implied by, the second.

According to rules (2)–(4), $F = V$ does not hold at the time it was initiated, while it holds at the time it was terminated.

The definitions of `initiatedAt` and `terminatedAt` are domain-specific. One common form of rule for `initiatedAt` is the following:

$$\begin{aligned}
\text{initiatedAt}(F = V, T) \leftarrow & \\
& \text{happensAt}(E, T), \\
& \text{Conditions}[T]
\end{aligned} \tag{5}$$

where $\text{Conditions}[T]$ is some set of further conditions referring to time-point T . `terminatedAt` rules are handled similarly. Note that in this EC formulation, `initiatedAt`($F = V, T$) does not necessarily imply that $F \neq V$ at T . Similarly, `terminatedAt`($F = V, T$) does not necessarily imply that $F = V$ at T . Suppose, for example, that $F = V$ is initiated at time-point 20 and terminated at time-point 30 and that there are no other time-points at which it is initiated or terminated. Then $F = V$ holds at all time-points T such that $20 < T \leq 30$. Suppose now that $F = V$ is initiated at time-points 10 and 20 and terminated at time-point 30 (and at no other time-points). Then $F = V$ holds at all T such that $10 < T \leq 30$. And suppose finally that $F = V$ is initiated at time-points 10 and 20 and terminated at time-points 25 and 30 (and at no other time-points). In that case $F = V$ holds at all T such that $10 < T \leq 25$. In Section 5 we illustrate the use of `initiatedAt` and `terminatedAt` rules for expressing LTA definitions.

4 Short-Term Activities

We use the first dataset of the CAVIAR project¹ to perform LTA recognition. This dataset includes 28 surveillance videos of a public space. The videos are staged — actors walk around, sit down, meet one another, leave objects behind, fight, and so on. Each video has been manually annotated by the CAVIAR team in order to provide the ground truth for both STA and LTA. For this set of experiments, the input to our recognition system is:

- (i) The STA *walking*, *running*, *active* (non-abrupt body movement in the same position) and *inactive* (standing still), together with their time-stamps, that is, the video frame in which that STA took place. The original CAVIAR dictionary does not include a STA for ‘abrupt motion’. Our preliminary experiments with this dataset showed that the absence of such a STA compromises the recognition accuracy of some LTA. ‘Abrupt motion’ is a form of STA that is detected by some state-of-the-art detection systems, for example (Kosmopoulos et al. 2008), but not by the CAVIAR systems. Accordingly, we modified the CAVIAR dataset by introducing a STA for ‘abrupt motion’: we manually edited the annotation of the CAVIAR videos by changing, when necessary, the label of a STA to *abrupt*. The STA *abrupt*, *walking*, *running*, *active* and *inactive* are mutually exclusive and represented by means of the `happensAt` predicate. For example, `happensAt(active(id6), 80)` expresses that *id₆* displayed ‘active’ bodily movement at time-point 80. STA are represented as instantaneous events in EC in order to use the `initiatedAt` and `terminatedAt` predicates to express the conditions in which these activities initiate and terminate a LTA.
- (ii) The coordinates of the tracked people and objects as pixel positions at each time-point, as well as their orientation. The coordinates are represented with the use of the `holdsAt` predicate. `holdsAt(coord(id2)=(14, 55), 10600)`, for example, expresses that the coordinates of *id₂* are (14, 55) at time-point 10600. Orientation is also encoded using the `holdsAt` predicate. For instance, `holdsAt(orientation(id2) = 120, 10600)` expresses that, in the two-dimensional projection of the video, the same person was forming a 120° angle with the x-axis at the same time-point. This type of information is necessary for computing the distance between two entities as well as the direction to which a person might be headed.
- (iii) The first and the last time a person or object is tracked (‘appears’ and ‘disappears’). This type of input is represented using the `happensAt` predicate. For example, `happensAt(appear(id10), 300)` expresses that *id₁₀* is first tracked at time-point 300.

Given such input, Crisp-EC recognises the following LTA: a person leaving an object, people meeting, moving together, or fighting. Long-term activities are represented as EC fluents. For instance, `holdsAt(moving(id1, id3) = true, 140)` states that *id₁* was moving together with *id₃* at time-point 140.

¹ <http://homepages.inf.ed.ac.uk/rbf/CAVIAR/>

LTA recognition is based on a manually developed knowledge base of LTA definitions expressed in terms of `initiatedAt` and `terminatedAt`. In the next section, we present example definition fragments of the LTA knowledge base. The complete code is available upon request.

5 Long-Term Activity Definitions

The ‘leaving an object’ activity is defined as follows:

$$\begin{aligned} \text{initiatedAt}(\text{leaving_object}(P, Obj) = \text{true}, T) \leftarrow \\ \text{happensAt}(\text{appear}(Obj), T), \\ \text{happensAt}(\text{inactive}(Obj), T), \\ \text{holdsAt}(\text{close}(P, Obj, 30) = \text{true}, T), \\ \text{holdsAt}(\text{person}(P) = \text{true}, T) \end{aligned} \quad (6)$$

$$\begin{aligned} \text{terminatedAt}(\text{leaving_object}(P, Obj) = \text{true}, T) \leftarrow \\ \text{happensAt}(\text{disappear}(Obj), T) \end{aligned} \quad (7)$$

In the CAVIAR videos an object carried by a person is not tracked — only the person that carries it is tracked. The object will be tracked (‘appear’) if and only if the person leaves it somewhere. Moreover, objects (as opposed to persons) can only exhibit the *inactive* STA. Accordingly, rule (6) expresses the conditions in which ‘leaving an object’ is recognised. The fluent recording this activity, *leaving_object*(*P*, *Obj*), becomes `true` at time *T* if *Obj* ‘appears’ at *T*, its STA at *T* is ‘inactive’, and there is a person *P* ‘close’ to *Obj* at *T*. The *close*(*ID1*, *ID2*, *Threshold*) fluent expresses that the distance between *ID1* and *ID2* is at most *Threshold* pixels. This fluent is defined as follows:

$$\begin{aligned} \text{holdsAt}(\text{close}(ID1, ID2, Threshold) = \text{true}, T) \leftarrow \\ \text{holdsAt}(\text{distance}(ID1, ID2) = Dist, T), \\ Dist < Threshold \end{aligned} \quad (8)$$

The distance between two tracked objects/people is computed as the Euclidean distance between their coordinates in the two-dimensional projection of the video — recall that the coordinates of each tracked entity are given as input to our system.

The 30 pixel distance threshold in rule (6) was determined from an empirical analysis of the CAVIAR dataset.

An object that is picked up by someone is no longer tracked (it ‘disappears’) which in turn terminates *leaving_object* — see rule (7).

In CAVIAR there is no explicit information that a tracked entity is a person or an inanimate object. Therefore, in our activity definitions we try to deduce whether a tracked entity is a person or an object given the detected STA. We defined the fluent *person*(*P*) to have value `true` if *P* was active, walking, running or moved

abruptly at some time-point since P ‘appeared’.

$$\begin{aligned}
& \text{initiatedAt}(\text{person}(P) = \text{true}, T) \leftarrow \\
& \quad \text{happensAt}(\text{active}(P), T) \\
& \text{initiatedAt}(\text{person}(P) = \text{true}, T) \leftarrow \\
& \quad \text{happensAt}(\text{walking}(P), T) \\
& \text{initiatedAt}(\text{person}(P) = \text{true}, T) \leftarrow \\
& \quad \text{happensAt}(\text{running}(P), T) \\
& \text{initiatedAt}(\text{person}(P) = \text{true}, T) \leftarrow \\
& \quad \text{happensAt}(\text{abrupt}(P), T) \\
& \text{terminatedAt}(\text{person}(P) = \text{true}, T) \leftarrow \\
& \quad \text{happensAt}(\text{disappear}(P), T)
\end{aligned} \tag{9}$$

The value of $\text{person}(P)$ is time-dependent because in CAVIAR the identifier P of a tracked entity that ‘disappears’ (is no longer tracked) at some point may be used later to refer to another entity that ‘appears’ (becomes tracked), and that other entity may not necessarily be a person. Note, finally, that rule (6) incorporates a (reasonable) simplifying assumption, that a person entity will never exhibit ‘inactive’ activity at the moment it first ‘appears’ (is tracked). If an entity is ‘inactive’ at the moment it ‘appears’ it can be assumed to be an object, as in the first two conditions of rule (6).

In a similar way, we may express the definitions of other LTA. The use of EC, in combination with the full power of logic programming, allows us to express LTA definitions including complex temporal, spatial or other constraints. Below we present fragments of the remaining LTA definitions.

meeting (of two persons P_1 and P_2) is recognised when two people ‘interact’: at least one of them is active or inactive, the other is not running or moving abruptly, and the distance between them is at most 25 pixels (all numeric constraints were determined from an empirical analysis of the dataset). In CAVIAR, this interaction phase can be seen as some form of greeting, such as a handshake. Rules (10) and (11) show the conditions in which *meeting* is initiated:

$$\begin{aligned}
& \text{initiatedAt}(\text{meeting}(P_1, P_2) = \text{true}, T) \leftarrow \\
& \quad \text{happensAt}(\text{inactive}(P_1), T), \\
& \quad \text{holdsAt}(\text{close}(P_1, P_2, 25) = \text{true}, T), \\
& \quad \text{holdsAt}(\text{person}(P_1) = \text{true}, T), \\
& \quad \text{holdsAt}(\text{person}(P_2) = \text{true}, T), \\
& \quad \text{not happensAt}(\text{running}(P_2), T), \\
& \quad \text{not happensAt}(\text{abrupt}(P_2), T), \\
& \quad \text{not happensAt}(\text{active}(P_2), T)
\end{aligned} \tag{10}$$

$$\begin{aligned}
& \text{initiatedAt}(\text{meeting}(P_1, P_2) = \text{true}, T) \leftarrow \\
& \quad \text{happensAt}(\text{active}(P_1), T), \\
& \quad \text{holdsAt}(\text{close}(P_1, P_2, 25) = \text{true}, T), \\
& \quad \text{holdsAt}(\text{person}(P_2) = \text{true}, T), \\
& \quad \text{not happensAt}(\text{running}(P_2), T), \\
& \quad \text{not happensAt}(\text{abrupt}(P_2), T)
\end{aligned} \tag{11}$$

meeting is terminated by a plethora of conditions, such as when one of the two people involved in the LTA starts running or ‘disappears’.

In CAVIAR *meeting* may have multiple initiations: two people may be interacting for several video frames. Similarly, *meeting* may have multiple terminations. This is in contrast to *leaving_object* where there is a single initiation and a single termination — for example, an object ‘appears’ once before it ‘disappears’. In general, there is no fixed relation between the number of initiations and terminations of a fluent — for example, a LTA may have multiple initiations and a single termination. In Section 3 we described how Crisp-EC computes the time-points in which a fluent with one or more initiations and terminations holds.

The activity *moving* was defined in order to recognise whether two people are walking along together:

$$\begin{aligned}
 \text{initiatedAt}(\text{moving}(P_1, P_2) = \text{true}, T) \leftarrow & \\
 \text{happensAt}(\text{walking}(P_1), T), & \\
 \text{happensAt}(\text{walking}(P_2), T), & \\
 \text{holdsAt}(\text{close}(P_1, P_2, 34) = \text{true}, T), & \quad (12) \\
 \text{holdsAt}(\text{orientation}(P_1) = Or_1, T), & \\
 \text{holdsAt}(\text{orientation}(P_2) = Or_2, T), & \\
 |Or_1 - Or_2| < 45 &
 \end{aligned}$$

In order to recognise *moving*, both people involved have to be *walking* while being close to each other. In addition, they have to be facing towards, more or less, the same direction (people walking in opposite directions are not assumed to be walking along together). This is accomplished by constraining their orientations so that they are, roughly, headed towards the same area while they are walking.

LTA — in contrast to STA — are not mutually exclusive. For example, *meeting* may overlap with *moving*: two people interact and then start *moving*, that is, they walk while being *close* to each other. In general, however, there is no fixed relationship between LTA.

moving is terminated when either person walks away from the other with respect to the predefined threshold of 34 pixels:

$$\begin{aligned}
 \text{terminatedAt}(\text{moving}(P_1, P_2) = \text{true}, T) \leftarrow & \\
 \text{happensAt}(\text{walking}(P_1), T), & \quad (13) \\
 \text{holdsAt}(\text{close}(P_1, P_2, 34) = \text{false}, T) &
 \end{aligned}$$

Other termination conditions for *moving* include either person running away from the other, as well as either person ‘disappearing’ from the scene.

The last definition concerns the *fighting* activity:

$$\begin{aligned}
 \text{initiatedAt}(\text{fighting}(P_1, P_2) = \text{true}, T) \leftarrow & \\
 \text{happensAt}(\text{abrupt}(P_1), T), & \quad (14) \\
 \text{holdsAt}(\text{close}(P_1, P_2, 44) = \text{true}, T), & \\
 \text{not happensAt}(\text{inactive}(P_2), T) &
 \end{aligned}$$

To recognise *fighting*, we require that both people are sufficiently close and at least one of them moves abruptly, while the other one is not *inactive*, indicating that

he ought to be participating in the fight somehow. *fighting* ceases to be recognised when either person involved in the activity starts walking away, running away, or exits the scene.

6 A Probabilistic Logic Programming Framework

In earlier work (Artikis, Sergot and Paliouras (2010)) we identified various types of uncertainty that exist in activity recognition, such as erroneous STA detection. To address this issue, we ported our EC dialect into ProbLog (Kimmig et al. 2011), a probabilistic extension of the logic programming language Prolog. ProbLog differs from Prolog in that it allows for probabilistic facts, which are facts of the form $p_i :: f_i$ where p_i is a real number in the range $[0, 1]$ and f_i is a Prolog fact. If f_i is not ground, then the probability p_i is applied to all possible groundings of f_i . Classic Prolog facts are silently given probability 1.

Probabilistic facts in a ProbLog program represent random variables. Furthermore, ProbLog makes an independence assumption on these variables. This means that a rule which is defined as a conjunction of n of these probabilistic facts has a probability equal to the product of the probabilities of these facts. When a predicate appears in the head of more than one rule then its probability is computed by calculating the probability of the implicit disjunction created by the multiple rules. For example, for a predicate p with two rules $p \leftarrow l_1$ and $p \leftarrow l_2, l_3$, the probability $P(p)$ is computed as follows:

$$\begin{aligned} P(p) &= P((p \leftarrow l_1) \vee (p \leftarrow l_2, l_3)) = \\ &= P(p \leftarrow l_1) + P(p \leftarrow l_2, l_3) - P((p \leftarrow l_1) \wedge (p \leftarrow l_2, l_3)) = \\ &= P(l_1) + P(l_2) \times P(l_3) - P(l_1) \times P(l_2) \times P(l_3) \end{aligned}$$

Given the independence assumption, any subprogram L has a probability equal to:

$$P(L) = \prod_{f_i \in L} p_i \cdot \prod_{f_i \notin L} (1 - p_i) \quad (15)$$

With the help of equation (15), one could compute the probability that a query q holds in a ProbLog program — *success probability* — by summing the probabilities of all subprograms that entail it:

$$P_s(q) = \sum_{L \models q} P(L) \quad (16)$$

Computing the success probability through equation (16), however, is computationally infeasible for large programs, since it involves summing through an exponential number of summands ($2^{|B_L|}$ different subprograms, where B_L is the Herbrand Base). By combining equations (15) and (16) and eliminating redundant terms, we end up with the following characterisation:

$$P_s(q) = P\left(\bigvee_{e \in \text{Proofs}(q)} \bigwedge_{f_i \in e} f_i \right) \quad (17)$$

That is, the task of computing the success probability of a query q is transformed into the task of computing the probability of the Disjunctive Normal Form (DNF)

formula of equation (17). Practically, equation (17) expresses that the success probability of query q is equal to the probability that at least one of its proofs is sampled. This, unfortunately, is not a question of straightforwardly transforming the probability of the DNF to a sum of products. If we were to translate equation (17) to a sum of products, we would assume that all different proofs (conjunctions) are disjoint, meaning that they represent mutually exclusive possible worlds, which does not hold in the general case. In order to make the proofs disjoint, one would have to enhance every conjunction with negative literals, in order to exclude worlds whose probability has already been computed in previous conjunctions of the DNF. This problem is known as the *disjoint-sum problem* and is known to be #P-hard (Valiant 1979).

ProbLog’s approach consists of using Binary Decision Diagrams (BDDs) (Bryant 1986) to compactly represent the DNF of equation (17). A BDD is a binary decision tree with redundant nodes removed and isomorphic subtrees merged. The BDD nodes represent the probabilistic facts of the ProbLog program. Every node has a ‘positive’ and ‘negative’ outward edge, leading to either a child node or the special ‘true’ or ‘false’ terminal nodes. The positive outward edge of the BDD node is labelled with the probability of the respective probabilistic fact and the negative edge is labelled with the complement of that probability. Positive and negative edges represent distinct decisions on inclusion of the relevant fact in the currently sampled possible world; a positive edge signifies that the fact represented by its parent node is included in the sample with the labelled probability, whereas a negative edge signifies that the fact is not included in the sample with the complement of the same probability. Therefore, by following a path from the root node to the ‘true’ terminal node, one could sample a conjunction of the DNF formula of equation (17). The ‘negative’ outward edges offer a compact representation of the negated literals required to enhance the DNF formula in order to make it represent a disjunction over disjoint conjunctions.

To summarise, ProbLog’s inference follows three general steps. The first step is to gather all proofs of the query q by scanning the Selective Linear Definite (SLD) tree of proofs and represent them as the DNF formula of equation (17). Afterwards, with the help of a built-in translation script, the DNF is translated to a BDD. Finally, the probability of this BDD is computed recursively, starting from the root node and assuming a probability of 1 for the ‘true’ terminal and 0 for the ‘false’ terminal.

With the help of BDDs, ProbLog inference is able to scale to queries containing thousands of different proofs (Kimmig et al. 2011). ProbLog’s efficiency was the driving force behind our decision to use this framework for activity recognition under uncertainty.

ProbLog has been fully integrated in the YAP Prolog system.² Further details, examples and code samples are available in (Kimmig et al. 2011) and on the ProbLog website.³

² <http://www.dcc.fc.up.pt/~vsc/Yap/>

³ <http://dtai.cs.kuleuven.be/problog/>

7 The Event Calculus in ProbLog

In this section we present the necessary transformations in order to make Crisp-EC ProbLog-compatible — the result is Prob-EC. We also explain the inference procedure of Prob-EC through two examples.

7.1 Transformation

To express our EC dialect in ProbLog we had to update our treatment of negation in order to allow for probabilistic atoms. To compute the complement of the probability of a probabilistic fact, ProbLog provides the built-in predicate `problog_not`. For any probabilistic fact $p_i :: f_i$, we have that:

$$P_s(\text{problog_not}(f_i)) = 1 - P_s(f_i) = 1 - p_i \quad (18)$$

`problog_not` can only be used on facts that are part of the knowledge base. For facts that do not belong to the knowledge base, such as a STA that is not part of an input stream, `problog_not` fails silently and reports a probability of 0. Consequently, translating `not` to `problog_not`, where facts are used in our EC axioms, would lead to undesirable behaviour. Consider, for example, rule (14) re-written by translating `not` to `problog_not`:

$$\begin{aligned} \text{initiatedAt}(\text{fighting}(P_1, P_2) = \text{true}, T) \leftarrow \\ \text{happensAt}(\text{abrupt}(P_1), T), \\ \text{holdsAt}(\text{close}(P_1, P_2, 44) = \text{true}, T), \\ \text{problog_not}(\text{happensAt}(\text{inactive}(P_2), T)) \end{aligned} \quad (14')$$

Furthermore, assume that id_1 and id_2 start *fighting* at some time-point t : they are both moving abruptly and their distance is less than 44 pixel positions. In this case, however, *fighting* will not be initiated at t as the final condition of rule (14') will not be satisfied: there are no *inactive* facts at t and therefore `problog_not` will fail. To overcome this issue, we defined the following predicate:

$$\text{negate}_1(\text{Fact}) \leftarrow \text{not Fact} \quad (19)$$

$$\text{negate}_1(\text{Fact}) \leftarrow \text{problog_not}(\text{Fact}) \quad (20)$$

With the use of `negate1` we produce a probability of 1 whenever the negated STA is not detected (see rule (19)), but also produce the complement of the probability of the STA whenever it is detected (see rule (20)). Note that in the latter case, that is, when a STA is detected with some probability, then ‘not STA’ fails.

Rule (14), for instance, can now be written as follows:

$$\begin{aligned} \text{initiatedAt}(\text{fighting}(P_1, P_2) = \text{true}, T) \leftarrow \\ \text{happensAt}(\text{abrupt}(P_1), T), \\ \text{holdsAt}(\text{close}(P_1, P_2, 44) = \text{true}, T), \\ \text{negate}_1(\text{happensAt}(\text{inactive}(P_2), T)) \end{aligned} \quad (14'')$$

With rule (14'') we are able to produce a probability of 1 whenever `happensAt(inactive(P2), T)` is not part of our input, as well as produce the complement of its probability whenever it is part of our input with a probability value

attached. The only case when `negate1` will produce a probability of 0 is when the probability value of its argument is 1. This is the desired behaviour.

The built-in predicate `problog_neg` is an extension of `problog_not` applicable to both probabilistic facts and derived atoms. For a derived atom r , we have:

$$P_s(\text{problog_neg}(r)) = 1 - P_s(r) \quad (21)$$

Similarly to `problog_not`, `problog_neg` cannot be used on goals that are not inferable from the knowledge base. Consider the following example:

$$\begin{aligned} \text{holdsAt}(F = V, T) \leftarrow \\ \text{initiatedAt}(F = V, T_s), \\ T_s < T, \\ \text{problog_neg}(\text{broken}(F = V, T_s, T)) \end{aligned} \quad (2')$$

This rule is produced by replacing `not` by `problog_neg` in rule (2). Assume that $F = V$ has been initiated with some probability p at $t-1$ and not `broken` in the meantime. In this case, `holdsAt`($F = V, t$) cannot be proved because the last condition of rule (2') fails. To overcome this issue, we defined the `negate2` predicate:

$$\text{negate}_2(\text{Goal}) \leftarrow \text{not Goal} \quad (22)$$

$$\text{negate}_2(\text{Goal}) \leftarrow \text{problog_neg}(\text{Goal}) \quad (23)$$

With the use of `negate2` we produce a probability of 1 whenever the negated goal is not inferable (see rule (22)), but also produce the complement of the probability of the goal whenever it is inferable (see rule (23)). Rule (2), for instance, can now be re-written as follows:

$$\begin{aligned} \text{holdsAt}(F = V, T) \leftarrow \\ \text{initiatedAt}(F = V, T_s), \\ T_s < T, \\ \text{negate}_2(\text{broken}(F = V, T_s, T)) \end{aligned} \quad (2'')$$

If $F = V$ is not `broken` then the last condition of rule (2'') will have probability 1 and, therefore, the probability that $F = V$ will depend entirely on the probability of its initiation conditions. If, however, $F = V$ is `broken` with some probability p , then the probability that $F = V$ will be equal to the product of the probability of its initiation conditions and $1-p$. This is the desired behaviour.

7.2 Inference

To illustrate the inference procedure of Prob-EC we use two LTA from CAVIAR, one with multiple initiations and terminations, and one with a single initiation and a single termination.

7.2.1 Multiple Initiations and Terminations

Suppose that `mike` and `sarah` are engaging in a ‘moving’ activity for a number of video frames — see Figure 1. The activity is first initiated at frame number 1, when

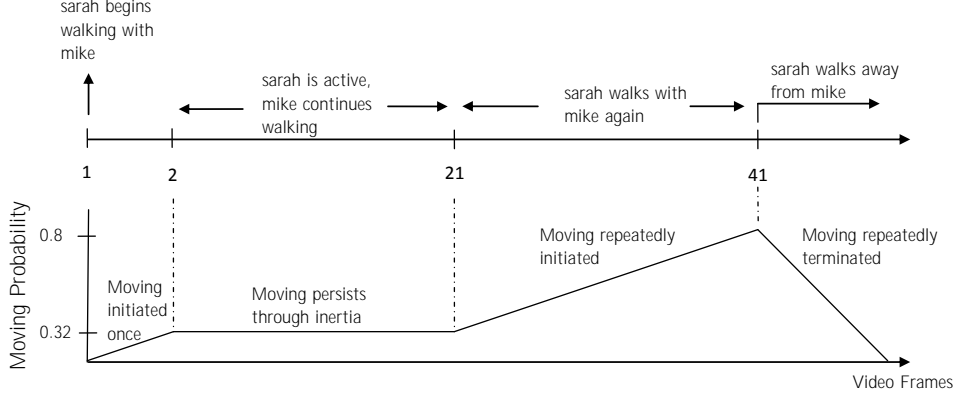


Fig. 1: Example of LTA probability fluctuation in the presence of multiple initiations and terminations.

both mike and sarah start *walking*. At frame 2 sarah stops *walking* (*walking* is required by rule (12) to initiate *moving*). She instead displays *active* body movement. Furthermore, mike continues *walking* but does not move far enough from her to terminate *moving*. At frame 21 sarah resumes *walking*, once again initiating *moving*. At frame 41 sarah continues *walking*, but mike is *inactive* and is left behind — sarah and mike are no longer close enough to each other, which triggers the termination condition (13) of *moving* (‘walk away’). For simplicity, let all information pertaining to orientation and coordinates be crisp (probability of 1). Moreover, assume that the STA *walking*, *active* and *inactive* have probabilities attached, as follows:

```

0.70 :: happensAt(walking(mike), 1)
0.46 :: happensAt(walking(sarah), 1)
0.73 :: happensAt(walking(mike), 2)
0.55 :: happensAt(active(sarah), 2)
...
0.69 :: happensAt(walking(mike), 21)
0.58 :: happensAt(walking(sarah), 21)
...
0.18 :: happensAt(inactive(mike), 41)
0.32 :: happensAt(walking(sarah), 41)
...

```

At frame 2, the query $\text{holdsAt}(\text{moving}(\text{mike}, \text{sarah}) = \text{true}, 2)$ has a probability equal to the probability of the initiation condition of frame 1, which, according to rule (12), and given that all coordinate and orientation-related information are crisply recognised, is the product of the probabilities that both mike and sarah are *walking*, that is, $0.70 \times 0.46 = 0.322$. This is visualised at the far left of Figure 1, where the LTA’s probability jumps from 0 to 0.322. From frame 2 to frame 20 no initiation or termination conditions are fired and the probability of *moving* remains unchanged. This occurs due to the law of inertia and is depicted graphically by the horizontal

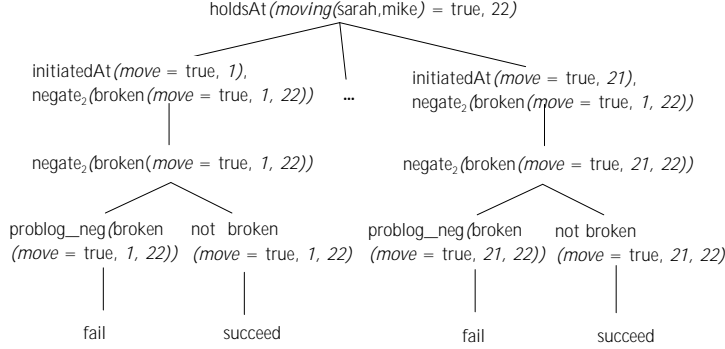


Fig. 2: SLD tree for query $\text{holdsAt}(\text{moving}(\text{mike}, \text{sarah}) = \text{true}, 22)$.

line between frames 2 and 20. At frame 21, *sarah* starts walking alongside *mike* again. Consequently, at frame 22, the query $\text{holdsAt}(\text{moving}(\text{mike}, \text{sarah}) = \text{true}, 22)$ has two initiation conditions to consider, one fired at frame 1 and one at frame 21. This occurs because rule (2'') searches over all time-points between time-point 0 and the current time-point for initiation conditions, finding both the condition fired at frame 1 and the one fired at frame 21.

As mentioned in the previous section, ProbLog computes the probability of a query by first scanning the entire SLD tree of the query. Figure 2 depicts a fragment of the SLD tree for the query $\text{holdsAt}(\text{moving}(\text{mike}, \text{sarah}) = \text{true}, 22)$. Then ProbLog represents these proofs as a DNF formula. In our case, the DNF is the following:

$$\underbrace{\text{initiatedAt}(\text{moving}(\text{sarah}, \text{mike}) = \text{true}, 1)}_{\text{init}_1} \vee \underbrace{\text{initiatedAt}(\text{moving}(\text{sarah}, \text{mike}) = \text{true}, 21)}_{\text{init}_{21}} \quad (24)$$

We have simplified the representation by omitting the two relevant ‘negate₂’ clauses since, as can be seen in Figure 2, they are both provable through negation as failure and therefore have a probability of 1. This occurs because no termination conditions for *moving* have been fired between frames 1 and 22.

Up to frame 22, there exist two initiation conditions for the *moving* LTA, init_1 and init_{21} (see formula (24)). In the general case, there may exist many more initiation conditions in the interval between the start of the video and the examined video frame. In addition, for every initiation condition, rule (2'') will check whether the LTA has been terminated by examining the interval between the initiation and the current video frame, repeating the process at the next video frame. This leads to numerous redundant computations. We overcame this problem by implementing an elementary caching technique according to which the probability of $\text{holdsAt}(F = V, T-1)$ is stored in memory and, therefore, $\text{holdsAt}(F = V, T)$ simply checks to see how the initiation or termination conditions (if any) fired at time-point $T-1$ affect this probability. This technique operates under the assumption that the activity recognition system receives the video frames in a temporally sorted manner — this assumption holds in CAVIAR.

The next step of ProbLog inference involves translating the DNF into a BDD. However, our example is simple enough to allow us to perform manual calculations, as there exist only two proofs for `holdsAt`, which are easy to disjoin. The probability of DNF formula (24) can be computed as the probability of a disjunction of two elements, as explained in Section 6:

$$\begin{aligned} P(\text{init}_1 \vee \text{init}_{21}) &= P(\text{init}_1) + P(\text{init}_{21}) - P(\text{init}_1 \wedge \text{init}_{21}) = \\ &= 0.70 \times 0.46 + 0.69 \times 0.58 - 0.7 \times 0.46 \times 0.69 \times 0.58 = 0.593 \end{aligned}$$

The probability that `mike` and `sarah` are *moving* at frame 22 has increased, owing to the presence of the additional initiation condition of frame 21. This is one of the characteristics of Prob-EC: the continuous presence of initiation conditions of a particular LTA causes an increase of the probability of the LTA. This behaviour is consistent with our intuition: given continuous indication that an activity has (possibly) occurred, we are more inclined to agree that it has indeed taken place, even if the confidence of every individual indication is low. For this reason, from frame 22 up to and including frame 41, the probability of *moving* increases, as is visible in Figure 1. In this example, at frame 41 the activity’s probability has escalated to around 0.8.

At frame 42, Prob-EC has to take into consideration the termination condition that was fired at frame 41. This termination condition, corresponding to rule (13), is also probabilistic: it bears the probability that `sarah` walked away from `mike`, which, according to rule (13) and the fact that `close` is crisply detected, is equal to the probability of the *walking* STA itself, which is 0.32. Therefore, when estimating the probability that, at frame 42, `mike` and `sarah` are still *moving* together, we have to incorporate the probability of all possible worlds in which `sarah` did not, in fact, walk away from `mike`. The probability of these worlds is computed by the use of `negate2` in rule (2’’) and is equal to $1 - 0.32 = 0.68$. Consequently, the probability that `mike` and `sarah` are still *moving* together at frame 42 is $0.8 \times 0.68 = 0.544$. Similarly to the steady probability increase given continuous initiation conditions, when faced with subsequent termination conditions, the probability of the LTA will steadily decrease. The slope of the descent (ascent) is defined by the probability of the termination (initiation) conditions involved. For this example, we assume that `sarah` keeps walking away from `mike` until the end of the video, causing the LTA’s probability to approximate 0, as shown at the far right of Figure 1.

Figure 3 shows the precise ProbLog output for our example. After an abrupt jump from 0 to 0.322, the probability remains stable between frames 2 and 21, indicating that for this period of time the LTA persists through the law of inertia. Between frames 22 and 41, the LTA’s probability monotonically increases, reflecting the repeated initiations of *moving* that occur during that time. After frame 41, it decreases, reflecting the repeated termination conditions occurring at the same time period. The dashed horizontal line at probability 0.5 represents the recognition threshold that we use to discern between LTA positives that we consider to be trustworthy enough — these are the LTA recognitions — and those that we do not.

Fig. 3: ProbLog output concerning $\text{holdsAt}(\text{moving}(\text{mike}, \text{sarah}) = \text{true}, T)$ for various video frames T .

The choice of a 0.5 threshold was made simply to provide a concrete illustration | other thresholds could have alternatively been used in this example.

7.2.2 Single Initiation and Termination

In this section we illustrate the behaviour of Prob-EC in the case of agents with a single initiation and a single termination. (Recall that there is no fixed relation between the number of initiations and terminations of an agent.) Assume that sarah is walking while simultaneously carrying a suitcase for 10 video frames. At frame 11, she leaves the suitcase on the floor and walks away from it. This causes the suitcase to 'appear' in the low-level tracking system, triggering the `leaving_object` initiation condition expressed by rule (6). Suppose that this initiation condition has a probability of 0.6. At frame 20, sarah picks up the suitcase, causing it to 'disappear' and triggering termination condition expressed by rule (7). Suppose that the termination condition also has a probability of 0.6. Figure 4 depicts the probability fluctuation of `leaving_object`. As can be seen in this figure, in the absence of any initiations after frame 11, the LTA persists entirely due to the law of inertia. The probability of this LTA is equal to the probability of the single initiation condition, that is, 0.6. Because this probability is above the chosen recognition threshold | 0.5 in this example | all frames taking place until sarah picks up the suitcase will be counted as recognitions. If the probability of the initiation condition was below the threshold then `leaving_object` would not have been recognised.

Note that we may transform an agent with a single initiation to a version of that agent with multiple initiations | consider the following formalisation:

$$\begin{aligned} &\text{initiatedAt}(\text{leaving_object_mi}(P; \text{Obj}) = \text{true}; T) \\ &\text{holdsAt}(\text{leaving_object}(P; \text{Obj}) = \text{true}; T) \end{aligned} \quad (25)$$

`leaving_object_mi` may have multiple initiations | it is initiated as long as `leaving_object` holds. While this transformation would have been beneficial for our experiments in the CAVIAR dataset, since we would be able to augment the prob-




Fig. 4: Example of LTA probability fluctuation in the presence of a single initiation and a single termination.

ability of *leaving_object_mi* through multiple initiations and eventually surpass the chosen recognition threshold, it introduces some subtle perils. Consider, for example, a scenario in which a *leaving_object* LTA is initiated with a small probability, such as 0.1, indicating that the sensor’s confidence about the STA appearing in rule (6) is low. Prob-EC will then compute that *leaving_object* holds with a probability of 0.1 up until the video frame, if any, where the object in question is picked up. These positives will be correctly discarded under most recognition thresholds given their small probabilities. *leaving_object_mi*, however, will continue to augment its probability, eventually surpassing the chosen threshold and thus producing a potentially large number of False Positives (FP). While it is true that such situations do not arise in CAVIAR, it is very likely that they take place in other activity recognition applications.

leaving_object has a single termination. In this example, the probability of the termination condition drops the probability of *leaving_object* below the chosen threshold of 0.5. In other examples, the absence of subsequent terminations may not allow the probability of a LTA to drop below the chosen recognition threshold, thus possibly resulting in false persistence.

8 Experimental Evaluation

8.1 CAVIAR without Artificial Noise

Our empirical analysis is based on the 28 surveillance videos of the CAVIAR dataset which contain, in total, 26419 video frames. These frames have been manually annotated by the CAVIAR team to provide the ground truth for STA and LTA (we performed very minor editing of the annotation in order to introduce a STA for abrupt motion). According to the manual annotation of the dataset, all STA are associated with a probability of 1. Table 2 shows the recognition results in terms of True Positives (TP), False Positives (FP), False Negatives (FN), Precision, Recall and F-measure. These results have been produced by computing queries of the form $\text{holdsAt}(LTA = \text{true}, T)$. Given that STA have probability of 1, Prob-EC has identical results to Crisp-EC.

Particularly notable in the results is the low Precision for the *leaving_object* LTA, owing to a substantial number of FP. This is due to the problematic annotation

LTA	TP	FP	FN	Precision	Recall	F-measure
meeting	3099	1910	525	0.619	0.855	0.718
moving	4008	2162	2264	0.650	0.639	0.644
fighting	531	97	729	0.421	0.845	0.562
leaving object	143	1539	55	0.085	0.722	0.152

Table 2: True Positives (TP), False Positives (FP), False Negatives (FN), Precision, Recall and F-measure on CAVIAR without artificial noise.

of CAVIAR with respect to this LTA. For example, in video 14, the object that is left at frame 946, triggering an initiation of *leaving_object*, is picked up (‘disappears’) at frame 1354. However, the relevant annotation mistakenly reports that *leaving_object* stops occurring at frame 996. We therefore end up with 358 FP which could have been avoided with a more consistent annotation of this video. Similarly, in the annotation of videos 17 and 18, a large number of annotated frames are missing. Video 16 includes a particularly interesting case of *leaving_object*. In this video, a person leaves a bag next to a chair, exits the scene, re-enters after a couple of seconds and picks up the bag. When the person re-enters, he is assigned a new identifier (this is common in CAVIAR). Various complications arise due to this. The original *leaving_object* activity is not terminated by our rules when the person in question ‘disappears’. This is deliberate on our part: we choose to terminate *leaving_object* when the object is picked up rather than when the person that leaves it ‘disappears’ from the sensor’s point of view. We thus emphasize on time-points in which a package might be unattended. The CAVIAR annotation, however, views the *leaving_object* LTA from a different perspective and thus assumes that the activity is terminated when the person ‘disappears’. This difference in perspective leaves us with a substantial number of FP, one for each frame that the person is not present in the scene. When the person re-enters the scene, CAVIAR provides the person with a new identifier and resumes the annotation of *leaving_object* with a $\langle new_person_id, same_object_id \rangle$ tuple. After a couple of frames, the person (described by *new_person_id*) picks up the object, terminating a *leaving_object(new_person_id, same_object_id)* LTA occurrence which Crisp-EC never initiated in the first place. Thus, in addition to a substantial number of FP, we also generate 55 FN (see Table 2) because Crisp-EC never recognises the new *leaving_object* activity.

8.2 CAVIAR with Artificial Noise

As mentioned above, according to the manual annotation of the CAVIAR dataset all STA are associated with a probability of 1. In real-world activity recognition applications it is unrealistic to assume that STA will be detected with certainty.


In order to experiment with a more realistic setting, we added artificial noise to the dataset, in the form of probabilities attached to the input facts — STA and related coordinate and orientation information. Our experimental procedure may be summarised as follows.

- We inject noise into CAVIAR:
 1. We add probabilities to STA. Toward this end, we use a Gamma distribution with a varying mean in order to represent different levels of noise.
 2. In addition to STA, we add probabilities to their associated coordinate and orientation fluents. Although we use the same Gamma distribution for this step, STA are not required to have the same probability as their associated coordinate and orientation fluents.
 3. On top of the above, we introduce spurious STA, that is, STA that are not part of the original CAVIAR dataset. We augment the frames in which there is a *walking* STA with another *walking* STA and related coordinate and orientation information about an entity that does not exist in CAVIAR. We use a uniform distribution to choose the frames that will be augmented with spurious facts. The probability of a spurious fact at frame t is $1-p$ where p is the probability of some CAVIAR *walking* STA at t (recall that p was computed by the Gamma distribution).

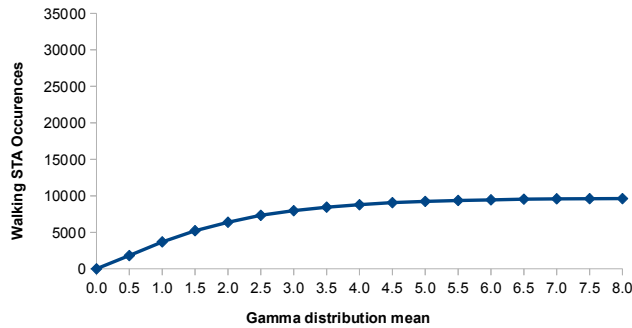
Thus we end up with three different noisy versions of CAVIAR. To facilitate the presentation that follows, we will call the three aforementioned approaches to noise injection ‘smooth’, ‘intermediate’ and ‘strong’ noise.

- We feed these data to Prob-EC and filter its output — which is a series of positives of the form $Prob :: \text{holdsAt}(LTA = \text{true}, T)$ — to keep only the positives with probability above a chosen threshold, indicating that we trust these positives to be accurate.
- We filter each noisy version of the dataset by erasing the facts with probability below the chosen threshold. We retain the facts with probability above the threshold, removing their probability values. Thus we assume that such facts have been tracked with certainty. For ‘smooth’ noise, this step means that we remove a certain amount of the CAVIAR STA, while for ‘intermediate’ noise we additionally remove coordinate and orientation fluents. Furthermore, for ‘strong’ noise we keep a certain amount of spurious information.
- We give these filtered versions of CAVIAR as input to Crisp-EC. With this step we aim to estimate the impact of environmental noise on Crisp-EC, by assuming that Crisp-EC can only reason on facts that have been tracked with relative certainty, that is, facts with probability above the chosen threshold.
- We compare the performance of Crisp-EC and Prob-EC.

We repeated the above experimental procedure 16 times, once for each Gamma distribution mean value between 0.5 and 8.0 inclusive, with a step of 0.5. Noise is added randomly; for example, a STA that is erased from Crisp-EC’s input for Gamma mean 6.0 might be present at the dataset produced for Gamma mean 6.5.



(a) (a) CAVIAR STA.



(b) (b) Spurious STA.

Fig. 5: Occurrences of *walking* STA with probability above 0.5 per Gamma mean value. Mean value 0.0 represents the CAVIAR dataset without artificial noise.

In our implementation, the higher the mean, the lower the probabilities attached to the CAVIAR input facts and the higher the probability of the spurious facts, indicating a higher amount of noise.

Figure 5(a) shows an aspect of noise injection. In this figure the number of occurrences of the *walking* STA with probability above the 0.5 threshold is plotted. Note that the numbers of *walking* occurrences shown in this figure do not include the spurious facts introduced in the ‘strong’ noise experiments. The amount of *walking* occurrences drops exponentially as we increase the level of noise, that is, the Gamma mean value. All other STA follow the same pattern. In the case of ‘intermediate noise’, the occurrences of coordinate and orientation fluents also drop exponentially.

Figure 5(b) shows the number of occurrences of the spurious *walking* STA — these are introduced in the ‘strong’ noise experiments — with probability above the 0.5 threshold. The amount of spurious facts increases exponentially as we increase the level of noise (Gamma mean value).

Noise injection may significantly reduce recognition accuracy. The sections that follow illustrate this. However, given that there are mistakes in the original dataset (that is, CAVIAR without artificial noise), it may occur that a FP in the original dataset becomes a True Negative (TN) in a noise-altered version of CAVIAR.

Similarly for FN. To demonstrate this with an example, assume that two people are *moving* along together on a pavement. Suddenly they have to step aside to allow a handicapped person full access to the pavement. This would fire the ‘walk away’ termination condition (see rule (13)) because the distance between the two people would exceed the pre-specified threshold. This does not mean that the people have stopped *moving* — they merely had to distance themselves momentarily. Firing — erroneously — the ‘walk away’ termination condition creates FN for some video frames. If, however, the *walking* STA gets a very low probability during the frames at which the people make way to the handicapped person, *moving* will not be terminated, which, in this special case, adds TP to the evaluation.




Fig. 6: Crisp-EC and Prob-EC F-measure per Gamma mean value under ‘smooth’ noise and a 0.5 threshold.

8.2.1 ‘Smooth’ Noise Experiments

Figure 6 compares the recognition accuracy of Crisp-EC and Prob-EC in terms of F-measure under ‘smooth’ noise using a 0.5 threshold. In all experiments we plot the F-measure per Gamma distribution mean averaged over 5 different runs, that is, each point in the diagram is the average of 5 different F-measure values. The vertical error bars display the standard deviations.

The recognition accuracy of Prob-EC is higher than that of Crisp-EC in *meeting*, *moving* and *fighting*. In *meeting* (see Figure 6(a)), the accuracy of Crisp-EC starts

falling from Gamma mean 5.5 onwards, because the *active* and *inactive* STA tend to be erased from its input, since they receive probability below 0.5 — the chosen threshold in the experiments presented in Figure 6 — when we generate noise. Prob-EC, on the other hand, is able to initiate *meeting* with a certain degree of probability (recall that rules (10) and (11) express the conditions in which *meeting* is initiated). After a number of frames, the repeated initiation of *meeting* leads to a *holdsAt* probability of 0.5 or above, and thus Prob-EC eventually recognises *meeting* (refer to Section 7.2.1 for a detailed explanation of this behaviour). Those initiation conditions occur very frequently in CAVIAR. Therefore, *meeting* ends up being recognised with a probability close to 1.

In the case of *moving* (see Figure 6(b)), it is the loss of multiple occurrences of *walking* (due to ‘smooth’ noise) that causes Crisp-EC to suffer from numerous FN (see rule (12) for the initiation condition of *moving*). Prob-EC, on the other hand, uses repeated initiation to eventually surpass the 0.5 recognition threshold. The accuracy of Crisp-EC in *moving* deteriorates earlier than that of *meeting*. This occurs because the initiation condition of *moving* bears two probabilistic conjuncts in its body, corresponding to two separate occurrences of the *walking* STA. Therefore, it is affected by noise twice. The initiation conditions of *meeting* may sometimes have two probabilistic conjuncts, but usually they have just one — the *negate*₁ conditions are usually crisp.

Similar to *moving* and *meeting*, Prob-EC fares better than Crisp-EC in *fighting* (see Figure 6(c)). For high levels of noise, that is, high Gamma mean values, Crisp-EC has trouble initiating *fighting* (see rule (14)), whereas Prob-EC uses repeated initiation to surpass the 0.5 recognition threshold. Prob-EC manages to surpass this threshold despite the relatively short duration of fights, compared to other LTA.

leaving_object is an interesting special case, owing to the fact that this LTA is recognised through a single initiation (see Section 7.2.2). In the CAVIAR videos, a person leaves an object after a few frames in which he is walking. In rare situations it is possible that either Crisp-EC or Prob-EC miss the recognition of *person* due to noise and, consequently, the recognition of *leaving_object*. Such cases did not arise in our ‘smooth’ noise experiments. In these experiments, Prob-EC and Crisp-EC are equally accurate (see Figure 6(d)). When *person* is recognised by Crisp-EC, it is recognised with a sufficiently high probability by Prob-EC and vice versa.

LTA in CAVIAR are usually terminated when an entity ‘disappears’ or when people walk away from each other. In the latter case, the probability of the LTA termination depends solely on that of *walking* (see, for example, rule (13)). In general, probabilistic terminations are similar to probabilistic initiations. When a termination condition is repeatedly fired with probabilities below 0.5 then Prob-EC eventually stops recognising the LTA, whereas Crisp-EC does not, producing FP.

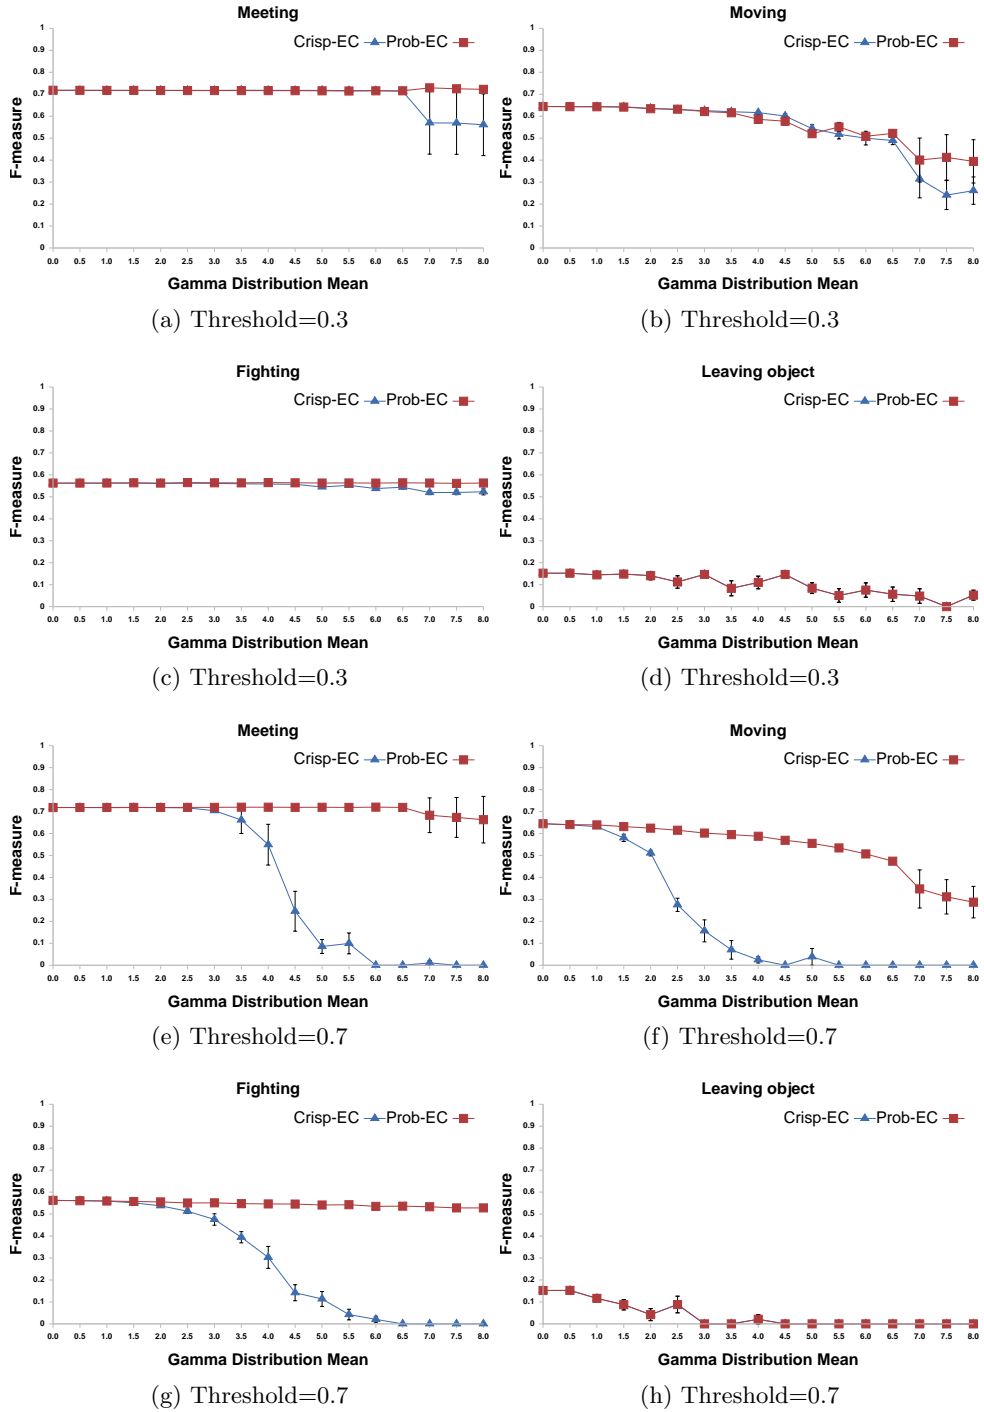


Fig. 7: Crisp-EC and Prob-EC F-measure per Gamma mean value under ‘smooth’ noise. In the first four figures the threshold is 0.3 while in the last four figures the threshold is 0.7.

Figure 7 compares the recognition accuracy of Crisp-EC and Prob-EC using 0.3 and 0.7 thresholds. This figure shows that Prob-EC is only slightly affected by the threshold change. This is due to the fact that the probabilities of most positives in Prob-EC are either very small (<0.3) or very high (>0.7). On the other hand, the recognition accuracy of Crisp-EC is highly affected by the choice of threshold: the lower the threshold the better the accuracy of Crisp-EC. This is expected given the nature of the ‘smooth’ noise experiments. As the threshold decreases, the input of Crisp-EC gets closer to the original, artificial noise-free dataset.




Fig. 8: Crisp-EC and Prob-EC F-measure per Gamma mean value under ‘intermediate’ noise and a 0.5 threshold.

8.2.2 ‘Intermediate’ Noise Experiments

Figure 8 compares the recognition accuracy of Crisp-EC and Prob-EC under ‘intermediate’ noise using a 0.5 threshold. Overall, attaching probabilities to the coordinates and orientation, in addition to STA, has reduced the accuracy of both Crisp-EC and Prob-EC. Compared to ‘smooth’ noise, the difference between Prob-EC and Crisp-EC in *meeting* is bigger (see Figure 8(a)). This occurs because, in addition to losing the *active* and *inactive* STA due to noise, Crisp-EC now has trouble proving that entities are *close*, because under ‘intermediate’ noise we also remove coordinate-related information, required to compute the distance between two entities. Thus, even in cases where the *active* or *inactive* STA are present and

indicate that the relevant frame might be an initiation condition for *meeting*, Crisp-EC is unable to prove that the two entities are *close* enough to initiate the LTA. Prob-EC, on the other hand, uses repeated initiation to eventually break the 0.5 barrier and produce recognitions.

Concerning *moving* (see Figure 8(b)), Crisp-EC suffers again from the loss of the *walking* STA. Under ‘intermediate noise’, this conclusion is more striking: after Gamma mean 2.5, Crisp-EC is unable to produce a single positive. This is because, in addition to the *walking* STA and associated coordinate fluents being erased from Crisp-EC’s input, orientation fluents are also candidates for removal. As a result, even in cases where both entities potentially involved in a *moving* activity are believed to be *walking* close to each other by the low-level tracking system, the absence of orientation information leads to an inability to initiate *moving*.

What is perhaps more interesting is that Prob-EC performs as bad as Crisp-EC in *moving* (see Figure 8(b)). This is because whenever faced with an initiation condition for *moving*, Prob-EC has to calculate the probability of this initiation condition as a product of 6 probabilities in total. (Recall from rule (8) that *close* is defined in terms of two coordinate input facts. It therefore contributes as the product of 2 probabilities.) Consequently, Prob-EC has trouble surpassing the 0.5 recognition threshold we require to trust its positives. Even in cases of near-certainty about some conditions of the input, ProbLog’s independence assumption leads to very low probability values. Consider, for example, two entities, id_1 and id_2 , whose STA (in our case, *walking*) and associated information (coordinates, orientation) are all tracked with a probability of 0.8. Whereas Crisp-EC will be able to initiate *moving*, since all the facts will make their appearance in the input, Prob-EC will produce an initiation condition probability of $(0.8)^6 = 0.262$. Consequently, we will not trust the probability of the relevant *holdsAt* query. Furthermore, due to the fact that each initiation condition has usually a very low probability, repeated initiation does not overcome the 0.5 threshold.

With respect to *fighting* (see Figure 8(c)), Prob-EC outperforms Crisp-EC for certain Gamma mean values. For high Gamma mean values, Prob-EC and Crisp-EC are equally inaccurate. In the case of ‘intermediate’ noise, there are at least three probabilistic conjuncts per initiation condition. Due to the fact that fights have a relatively short duration (much shorter than *meeting*, for example), there are not enough initiations to raise Prob-EC’s probability above the 0.5 threshold in high Gamma mean values.

Regarding *leaving_object* (see Figure 8(d)), Crisp-EC seems to fair slightly better than Prob-EC for low Gamma mean values. Under the ‘intermediate’ noise assumption, Prob-EC has to consider more probabilistic conjuncts per initiation condition, due to the presence of *close*. As a result, and given the single initiation of *leaving_object*, Prob-EC tends to produce probabilities below 0.5 for this LTA, even in cases where the STA probabilities themselves might be above 0.5 and therefore sufficient to allow Crisp-EC to recognise *leaving_object*. For higher Gamma mean values, Prob-EC and Crisp-EC are equally inaccurate, because the data required by Crisp-EC to initiate *leaving_object*, whether it is the *person* fluent, the *inactive* STA or the coordinate fluents of the entities involved, are missing.




Fig. 9: Crisp-EC and Prob-EC F-measure per Gamma mean value under ‘intermediate’ noise. In the first four figures the threshold is 0.3 while in the last four figures the threshold is 0.7.

Figure 9 compares the recognition accuracy of Crisp-EC and Prob-EC using 0.3 and 0.7 thresholds. As in the case of the ‘smooth’ noise experiments, Crisp-EC is affected more than Prob-EC by the threshold change — the accuracy of Crisp-EC improves as the threshold decreases (the lower the threshold, the closer the input of Crisp-EC to the original, artificial noise-free dataset).

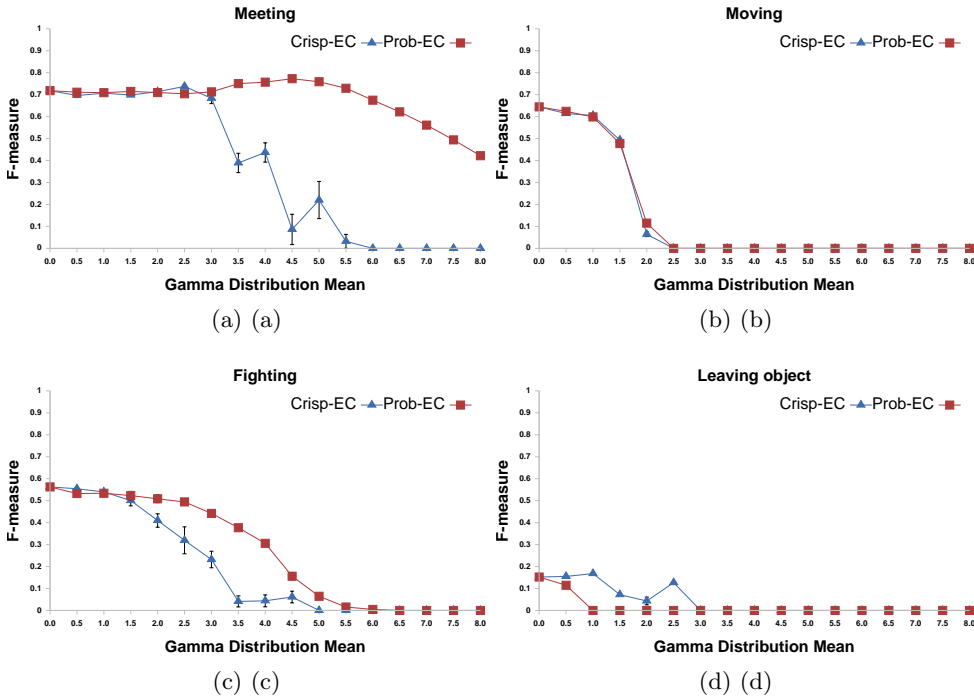


Fig. 10: Crisp-EC and Prob-EC F-measure per Gamma mean value under ‘strong’ noise and a 0.5 threshold.

8.2.3 ‘Strong’ Noise Experiments

Figure 10 compares the recognition accuracy of Crisp-EC and Prob-EC under ‘strong’ noise using a 0.5 threshold. As can be seen from this figure, the recognition accuracy under ‘strong’ noise is very similar to that of ‘intermediate’ noise in the 0.5 threshold. For Prob-EC, we need a series of spurious information to surpass the chosen threshold — the introduction of spurious STA in the ‘strong’ noise experiments is not that systematic and does not create problems to Prob-EC. Crisp-EC is very slightly affected by the introduction of spurious STA: it now has a number of additional FP, mainly in the case of *moving*. Even a single initiation condition fired by a spurious *walking* STA and some other CAVIAR STA may be enough in Crisp-EC to create a series of FP. However, in low Gamma values most of the spurious STA have low probabilities — below the 0.5 threshold — and therefore are not part of the input of Crisp-EC. Furthermore, in high Gamma values the

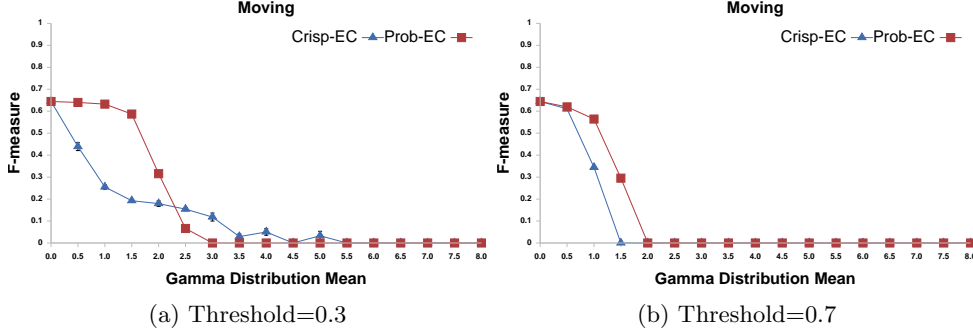


Fig. 11: Crisp-EC and Prob-EC F-measure per Gamma mean value under ‘strong’ noise in *moving*.

spurious STA cannot initiate a LTA because several of the CAVIAR STA tend to be deleted from the input of Crisp-EC as they have low probabilities.

Figure 11 compares the accuracy of Crisp-EC and Prob-EC in *moving* using thresholds of 0.3 and 0.7 (the diagrams for the remaining LTA are omitted as they are similar to those concerning ‘intermediate’ noise). Crisp-EC, as in the previous experimental settings, is more significantly affected than Prob-EC by the threshold change. Unlike the last two experimental settings, reducing the threshold causes problems to Crisp-EC with respect to *moving*. With a lower threshold, Crisp-EC loses less CAVIAR STA, but at the same time keeps more spurious facts, creating many more FP. Prob-EC, even in the low threshold of 0.3, is almost unaffected by the spurious facts. In the 0.7 threshold Crisp-EC has a very small number of spurious facts in its input and, therefore, its accuracy is not compromised by these facts.

In addition to *walking* STA, we could have added other spurious facts such *abrupt*, *active*, or *inactive* and repeated the experiments. The results would have been similar to those presented above — for example, spurious *abrupt* STA would have compromised the accuracy of Crisp-EC in *fighting*, especially in low thresholds.

9 Summary and Future Work

We presented Prob-EC, an Event Calculus dialect for probabilistic reasoning. Prob-EC is the first Event Calculus dialect able to deal with uncertainty in the input STA. Moreover, this is the first approach that thoroughly evaluates the Event Calculus in a probabilistic framework. Our experimental evaluation on a benchmark activity recognition dataset showed that Prob-EC outperforms Crisp-EC when:

- a LTA has multiple initiations, and
- the LTA depends on a small number of probabilistic conjuncts.

When a LTA depends on a large number of probabilistic conjuncts, Prob-EC is significantly affected and its performance is close to that of Crisp-EC. ProbLog makes an independence assumption about input facts and thus the product of the

probabilities of many probabilistic conjuncts may be very small, even if the probability of each individual conjunct is high. The greater the number of probabilistic conjuncts, the more initiations Prob-EC requires to surpass the chosen recognition threshold.

Prob-EC is more resilient to spurious facts than Crisp-EC. Even a single such fact may create a series of false positives (FP) in Crisp-EC, whereas this type of noise must be much more systematic to affect Prob-EC: to impinge the accuracy of the latter we need a series of spurious facts in order to surpass the recognition threshold.

In the case of a single initiation, there are situations in which Prob-EC may fare significantly better than Crisp-EC and vice versa, but these did not arise in our experiments. In general, Prob-EC is expected to have similar performance to Crisp-EC on LTA with a single initiation in the case of a small number of probabilistic conjuncts, while Crisp-EC is likely to perform better in the case of a large number of such conjuncts.

The experimental results concerning one or more terminations are similar to those regarding one or more initiations. For example, Prob-EC outperforms Crisp-EC in the case of LTA with multiple terminations that depend on a small number of probabilistic conjuncts. The repeated terminations allow Prob-EC to stop recognising a LTA when Crisp-EC continues the recognition producing FP.

There are several directions for further work. First, we intend to experiment with additional types of noise and additional datasets. Second, we aim to incorporate advanced caching techniques in Prob-EC, such as those presented in (Artikis, Sergot and Paliouras (2012)), in order to make it suitable for run-time activity recognition. Finally, we aim to address the issue of imprecise LTA definitions — we intend to provide a unified framework that will be able to deal with both STA detection probabilities and imperfect LTA definitions.

Acknowledgements

This work was supported by the EU PRONTO Project (FP7-ICT 231738). We are very grateful to the ProbLog developing team whose feedback contributed greatly to our grasp of the language. The authors themselves, however, are solely responsible for any misunderstanding about its use. We have also benefited from discussions with Marek Sergot on the Event Calculus.

References

- ALLEN, J. 1983. Maintaining knowledge about temporal intervals. *Communications of the ACM* 26, 11, 832–843.
- ARTIKIS, A., SERGOT, M., AND PALIOURAS, G. 2010. A logic programming approach to activity recognition. In *Proceedings of ACM Workshop on Events in Multimedia*.
- ARTIKIS, A., SERGOT, M., AND PALIOURAS, G. 2012. Run-time composite event recognition. In *Proceedings of International Conference on Distributed Event-Based Systems (DEBS)*. ACM, 69–80.

- ARTIKIS, A., SKARLATIDIS, A., PORTET, F., AND PALIOURAS, G. 2012. Logic-based event recognition. *Knowledge Engineering Review* 27, 4, 469–506.
- BISWAS, R., THRUN, S., AND FUJIMURA, K. 2007. Recognizing activities with multiple cues. In *Proceedings of Workshop on Human Motion*, A. M. Elgammal, B. Rosenhahn, and R. Klette, Eds. LNCS, vol. 4814. Springer, 255–270.
- BRAND, M., OLIVER, N., AND PENTLAND, A. 1997. Coupled Hidden Markov Models for complex action recognition. In *Proceedings of International Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, 994–999.
- BRENDEL, W., FERN, A., AND TODOROVIC, S. 2011. Probabilistic event logic for interval-based event recognition. In *Proceedings of International Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 3329–3336.
- BRUYNOOGHE, M., MANTADELIS, T., KIMMIG, A., GUTMANN, B., VENNEKENS, J., JANSSENS, G., AND RAEDT, L. D. 2010. ProbLog technology for inference in a probabilistic first order logic. In *Proceedings of European Conference on Artificial Intelligence (ECAI)*. 719–724.
- BRYANT, R. 1986. Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers* 35(8), 677–691.
- CUGOLA, G. AND MARGARA, A. 2012. Processing flows of information: From data stream to complex event processing. *ACM Computing Surveys* 44, 3, 15.
- DOUSSON, C. AND MAIGAT, P. L. 2007. Chronicle recognition improvement using temporal focusing and hierarchisation. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*. 324–329.
- FIERENS, D., DEN BROECK, G. V., THON, I., GUTMANN, B., AND RAEDT, L. D. 2011. Inference in probabilistic logic programs using weighted CNF's. In *Proceedings of International Conference on Uncertainty in Artificial Intelligence (UAI)*. 211–220.
- GIBSON, J. 1979. The ecological approach to visual perception.
- GINSBERG, M. L. 1990. Bilattices and modal operators. *Journal of Logic and Computation* 1, 1–41.
- GONG, S. AND XIANG, T. 2003. Recognition of group activities using dynamic probabilistic networks. In *Proceedings of International Conference on Computer Vision*. IEEE, 742–749.
- HAKHEEM, A. AND SHAH, M. 2007. Learning, detection and representation of multi-agent events in videos. *Artificial Intelligence* 171, 8–9, 586–605.
- HELAOUI, R., NIEPERT, M., AND STUCKENSCHMIDT, H. 2011. Recognizing interleaved and concurrent activities: A statistical-relational approach. In *Proceedings of International Conference on Pervasive Computing and Communications*. IEEE, 1–9.
- HONGENG, S. AND NEVATIA, R. 2003. Large-scale event detection using semi-Hidden Markov Models. In *Proceedings of International Conference on Computer Vision*. IEEE, 1455–1462.
- KEMBHAVI, A., YEH, T., AND DAVIS, L. S. 2010. Why did the person cross the road (there)? scene understanding using probabilistic logic models and common sense reasoning. In *Proceedings of European Conference on Computer Vision (ECCV)*. 693–706.
- KIMMIG, A., DEMOEN, B., RAEDT, L. D., COSTA, V. S., AND ROCHA, R. 2011. On the implementation of the probabilistic logic programming language ProbLog. *Theory and Practice of Logic Programming* 11, 235–262.
- KOSMOPOULOS, D., ANTONAKAKI, P., VALASOULIS, K., KESIDIS, A., AND PERANTONIS, S. 2008. Human behaviour classification using multiple views. In *Proceedings of Hellenic Conference on Artificial Intelligence*. Vol. 5138. Springer.
- KOWALSKI, R. AND SERGOT, M. 1986. A logic-based calculus of events. *New Generation Computing* 4, 1, 67–96.

- LAFFERTY, J. D., MCCALLUM, A., AND PEREIRA, F. C. N. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of International Conference on Machine Learning (ICML)*, C. E. Brodley and A. P. Danyluk, Eds. Morgan Kaufmann, 282–289.
- LIAO, L., FOX, D., AND KAUTZ, H. 2007. Hierarchical conditional random fields for GPS-based activity recognition. *Robotics Research*, 487–506.
- LUCKHAM, D. 2002. *The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems*. Addison-Wesley.
- MOLDOVAN, B., MORENO, P., VAN OTTERLO, M., SANTOS-VICTOR, J., AND DE RAEDT, L. 2012. Learning relational affordance models for robots in multi-object manipulation tasks. In *Proceedings of International Conference on Robotics and Automation (ICRA)*. 4373–4378.
- MORARIU, V. I. AND DAVIS, L. S. 2011. Multi-agent event recognition in structured scenarios. In *Proceedings of International Conference on Computer Vision and Pattern Recognition (CVPR)*. 3289–3296.
- MURPHY, K. 2002. Dynamic bayesian networks: representation, inference and learning. Ph.D. thesis, University of California.
- NATARAJAN, P. AND NEVATIA, R. 2007. Hierarchical multi-channel semi Hidden Markov Models. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*. 2562–2567.
- RABINER, L. AND JUANG, B. 1986. An introduction to Hidden Markov Models. *ASSP Magazine* 3, 1, 4–16.
- RICHARDSON, M. AND DOMINGOS, P. 2006. Markov logic networks. *Machine Learning* 62, 1-2, 107–136.
- SADILEK, A. AND KAUTZ, H. 2012. Location-based reasoning about complex multi-agent behavior. *Journal of Artificial Intelligence Research* 43, 87–133.
- SELMAN, J., AMER, M., FERN, A., AND TODOROVIC, S. 2011. PEL-CNF: Probabilistic event logic conjunctive normal form for video interpretation. In *Computer Vision Workshops*. IEEE, 680–687.
- SHET, V., HARWOOD, D., AND DAVIS, L. 2005. VidMAP: video monitoring of activity with Prolog. In *Proceedings of International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. IEEE, 224–229.
- SHET, V., NEUMANN, J., RAMESH, V., AND DAVIS, L. 2007. Bilattice-based logical reasoning for human detection. In *Proceedings of International Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 1–8.
- SISKIND, J. 2001. Grounding the lexical semantics of verbs in visual perception using force dynamics and event logic. *Journal of Artificial Intelligence Research* 15, 31–90.
- SKARLATIDIS, A., PALIOURAS, G., VOUIROS, G. A., AND ARTIKIS, A. 2011. Probabilistic event calculus based on markov logic networks. In *RuleML America*. 155–170.
- TRAN, S. D. AND DAVIS, L. S. 2008. Event modeling and recognition using markov logic networks. In *Proceedings of European Conference on Computer Vision (ECCV)*. 610–623.
- VAIL, D., VELOSO, M., AND LAFFERTY, J. 2007. Conditional random fields for activity recognition. In *Proceedings of International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. ACM, 1–8.
- VALIANT, L. G. 1979. The complexity of enumeration and reliability problems. *SIAM Journal on Computing* 8, 410–421.
- WANG, J. AND DOMINGOS, P. 2008. Hybrid Markov logic networks. In *Proceedings of Conference on Artificial Intelligence (AAAI)*. 1106–1111.

- WU, T., LIAN, C., AND HSU, J. 2007. Joint recognition of multiple concurrent activities using factorial conditional random fields. In *Proceedings of AAAI Workshop on Plan, Activity, and Intent Recognition*.