Interval-based Activity Recognition

Evangelos Makris[1], Alexander Artikis[2,1], and Georgios Paliouras[1]

Institute of Informatics and Telecommunications, NCSR "Demokritos"[1]
Department of Maritime Studies, University of Piraeus[2]
{vmakris,a.artikis,paliourg}@iit.demokritos.gr

Abstract. Activity recognition refers to the detection of temporal combinations of 'low-level' or 'short-term' activities on sensor data. Various types of uncertainty exist in activity recognition systems and this often leads to erroneous detection. Typically, the frameworks aiming to handle uncertainty compute the probability of the occurrence of activities at each time-point. We extend this approach by defining the probability of a maximal interval and the credibility rate for such intervals.

1 Introduction

In activity recognition, multiple sources provide spatial and temporal data that can be used to detect various types of human behaviour. The input data are *short-term* activities (STA), such as 'walking, 'running, 'active' and 'inactive', indicating that a person is walking, running, moving his arms while in the same position, and so on. The output is a set of *long-term* activities (LTA), which are temporal combinations of STA. Examples are 'fighting', 'meeting', 'moving', etc. When a rule that consists of temporal constraints on a set of STA is satisfied, it leads to the recognition of LTA. Uncertainty is inherent in activity recognition. For example, STA, typically detected by visual information processing tools operating on video feeds, often have probabilities attached to them by low-level classifiers, serving as a confidence estimates. In earlier work, we presented an activity recognition system based on a probabilistic version of the Event Calculus, hereafter Prob-EC, that computes the probability of an LTA at each time-point [12]. We extend this approach by defining the probability of a maximal interval and the credibility rate for such intervals. In contrast to time-point-based activity recognition, our proposed method is robust to noisy LTA probability fluctuations.

2 Background

2.1 Event Calculus

We restrict attention to a simple version of the Event Calculus where the time model is linear and includes integer time-points. Variables start with an upper-case letter, while predicates and constants start with a lower-case letter. 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. The domain-independent axioms are presented below:

$$\begin{aligned}
&\mathsf{holdsAt}(F = V,\ T) \leftarrow \\
&\quad \mathsf{initiatedAt}(F = V,\ T_s),\ T_s < T, \\
&\quad \mathsf{not\ broken}(F = V,\ T_s,\ T).
\end{aligned} \tag{1}$$

$$\begin{aligned}
&\mathsf{broken}(F = V,\ T_s,\ T) \leftarrow \\
&\quad \mathsf{terminatedAt}(F = V,\ T_f),\ T_s < T_f < T.
\end{aligned} \tag{2}$$

$$\begin{aligned}
&\mathsf{broken}(F = V,\ T_s,\ T) \leftarrow \\
&\quad \mathsf{initiatedAt}(F = V',\ T_f),\ V \neq V',\ T_s < T_f < T.
\end{aligned} \tag{3}$$

According to axiom (1), $F = V$ holds at some time-point T if it has been initiated by an event previously and has not been 'broken' in the meantime. This expresses the law of inertia. $F = V$ is 'broken' in (T_s, T) if it is terminated (see axiom (2)) or $F = V'$ is initiated, for some $V' \neq V$ (see axiom (3)). The definitions of initiatedAt and terminatedAt are domain-specific. Consider, for example, the (partial) definition of *moving* from the domain of activity recognition:

$$\begin{aligned}
&\mathsf{initiatedAt}(moving(P_1, P_2) = \mathsf{true},\ T) \leftarrow \\
&\quad \mathsf{happensAt}(walking(P_1),\ T), \\
&\quad \mathsf{happensAt}(walking(P_2),\ T), \\
&\quad \mathsf{holdsAt}(close(P_1, P_2) = \mathsf{true},\ T), \\
&\quad \mathsf{holdsAt}(similarOrientation(P_1, P_2) = \mathsf{true},\ T).
\end{aligned} \tag{4}$$

$$\begin{aligned}
&\mathsf{terminatedAt}(moving(P_1, P_2) = \mathsf{true},\ T) \leftarrow \\
&\quad \mathsf{happensAt}(walking(P_1),\ T), \\
&\quad \mathsf{holdsAt}(close(P_1, P_2) = \mathsf{false},\ T).
\end{aligned} \tag{5}$$

moving is a long-term activity (LTA) expressed as a Boolean fluent, and defined in terms of a set of short-term activities (STA) expressed as instantaneous events, and contextual information detected on video content. *walking*, *running*, *active* and *inactive* are mutually exclusive STA detected on video frames. Each such STA is accompanied by the coordinates and orientation of the tracked entity in question. These are the input of the activity recognition system. $close(P_1, P_2)$ is true when the distance between the tracked entities P_1 and P_2 is smaller than some pre-defined threshold of pixel positions. Similarly, $similarOrientation(P_1, P_2)$ is true when the difference in orientation of P_1 and P_2 is less than 45 degrees. According to rule (4), $moving(P_1, P_2) = \mathsf{true}$ is said to be initiated when both P_1 and P_2 are walking, they are close to each other and have a similar orientation. Furthermore, $moving(P_1, P_2) = \mathsf{true}$ is said to be terminated when the two tracked persons walk away from each other (see rule (5)). The remaining terminating conditions are defined in a similar manner [12].

Note that $\mathsf{initiatedAt}(F = V, T)$ does not necessarily imply that $F \neq V$ at T. Similarly, $\mathsf{terminatedAt}(F = V, T)$ does not necessarily imply that $F = V$ at T [3]. Suppose 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 $10 < T \leq 25$.

2.2 Point-based Probabilistic Event Calculus

Prob-EC [12] is a probabilistic version of the Event Calculus implemented in ProbLog [7]. The aim of Prob-EC is to compute the probabilities of $\mathsf{holdsAt}(F = V, T)$, i.e. the truth value of $F = V$ at time-point T. A Prob-EC programme consists of probabilistic facts, the domain-indepedent rules of the Event Calculus (see rules (1)–(3)), as well as domain-specific rules (such as rules (4) and (5)). Probabilistic facts are defined as $p :: f$, meaning that f holds as true with probability p in each of its groundings. All these facts represent independent random variables. The marginal probability of the head holding is therefore the product of the probabilities of the facts holding. This way, Prob-EC deals with uncertainty in the input data. The probability of $\mathsf{holdsAt}(LTA = \mathsf{true}, T)$ is equal to the probability of the disjunction of the initiation conditions of $LTA = \mathsf{true}$ before T, assuming that $LTA = \mathsf{true}$ has not been 'broken' in the meantime. Hence, multiple initiations of $LTA = \mathsf{true}$ increase its probability. Moreover, if $LTA = \mathsf{true}$ is 'broken' with probability p_1, then the probability of $LTA = \mathsf{true}$ becomes equal to the product of the probability of the disjunction of initiations and $1 - p_1$ (see axiom (1)). Therefore, the higher the probability p_1 the more significant the decrease of the probability of $LTA = \mathsf{true}$. Furthermore, consecutive terminations decrease further the probability of $LTA = \mathsf{true}$ [12].

3 Probabilistic Maximal Interval Estimation

An instantaneous indication of an activity, by means of $\mathsf{holdsAt}$, for example, may lead to erroneous detection, which may be due to the unreliability of the sensors or due to the inaccuracy of the recognition patterns. Towards this, we propose a *Probabilistic Interval−based Event Calculus* (*PIEC*). Figure 1 shows a high-level description of the inference procedure. First, we use Prob-EC, as described in the previous section, to compute the probabilities of LTA at each time-point given the probabilistic 'short-term' activities (STA). The recognition is based on domain-specific rules of initiation and termination, such as rules (4)–(5). The next phase consists of the interval-based activity recognition. With respect to a probability threshold, PIEC computes all 'probabilistic maximal intervals', i.e. the maximal intervals within which an activity is likely to hold.

Definition 1. *The probability of interval* $I_{LTA} = [i, j]$ *of LTA with* $length(I_{LTA}) = j - i + 1$ *time-points is defined as*

$$P(I_{LTA}) = \frac{\sum_{k=i}^{j} P(\mathit{holdsAt}(LTA, k))}{length(I_{LTA})}.$$

In other words, the probability of an interval is equal to the average of the probabilities at the time-points that it contains.

A key concept of PIEC is that of probabilistic maximal interval:

Definition 2. *A probabilistic maximal interval* $I_{LTA} = [i, j]$ *of LTA is an interval such that, given some threshold* $\mathcal{T} \in [0, 1]$, $P(I_{LTA}) \geq \mathcal{T}$, *and there is no other interval* I'_{LTA} *such that* $P(I'_{LTA}) \geq \mathcal{T}$ *and* I_{LTA} *is a sub-interval of* I'_{LTA}.

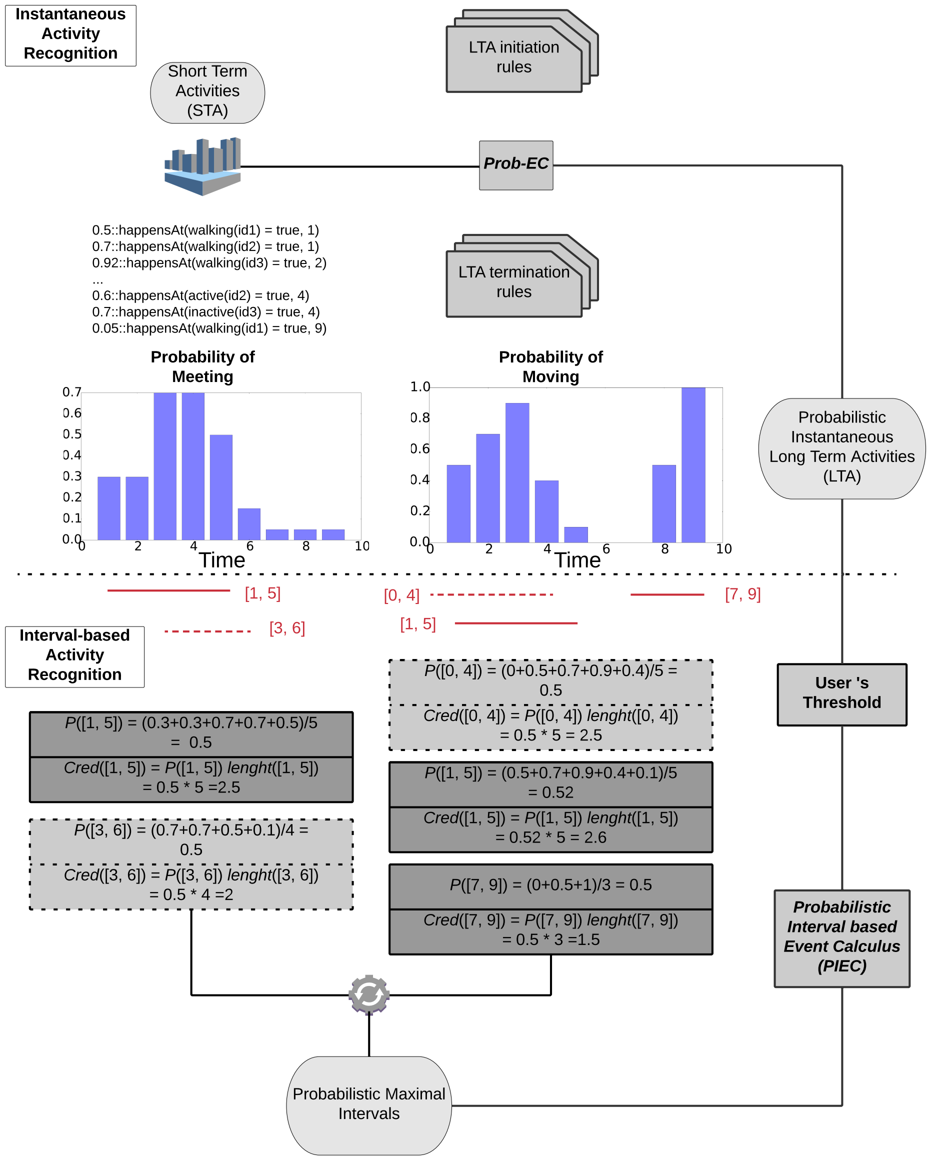

Fig. 1: Interval-based activity recognition. First (top figure, above the dashed line), Prob-EC computes the instantaneous probabilities of LTA, such as *meeting* and *moving*, given the probabilistic STA, such as 'walking', 'active' and 'inactive'. Second (bottom figure, below the dashed line), PIEC computes the 'probabilistic maximal intervals' of LTA, assuming a probability threshold (0.5 in this example). These intervals are depicted by the red (dashed) lines below the instantaneous probability evolution diagrams. The computation of the probability and 'credibility' of each such interval is presented in the boxes below the red lines.

A consequence of the definition of a probabilistic maximal interval is that such intervals may be overlapping. Two examples are shown in Figure 1—see the overlapping red lines under the instantaneous probability evolution diagrams of *meeting* and *moving*. From a set of overlapping probabilistic maximal intervals, we keep only one, using interval 'credibility', defined as the product of interval length and probability:

$$Cred(I_{LTA}) = length(I_{LTA}) \cdot P(I_{LTA}) = \sum_{k} P(\mathsf{holdsAt}(LTA, k)), \qquad (6)$$

where k are the time-points of the interval I_{LTA}. Hence, for each set of overlapping probabilistic maximal intervals $S = \{I_1, I_2, \ldots, I_k\}$, we select the one with the highest credibility, i.e. we select I_{LTA} with $Cred(I_{LTA}) = \max_i(Cred(I_i))$ for $i = 1, \ldots, k$. In Figure 1 the credible intervals are depicted by the solid red lines. Equation (6) ensures that we keep an interval which is as likely and long as possible. Nevertheless, this is just one of the many ways of picking between overlapping probabilistic maximal intervals.

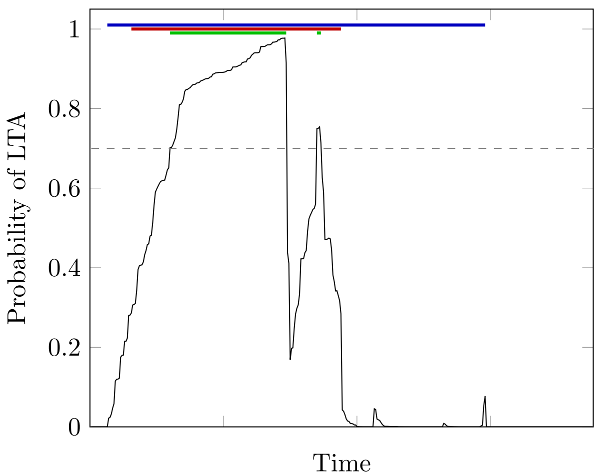

Fig. 2: Probabilistic activity recognition. The black line represents the LTA probability evolution computed by Prob-EC. The green horizontal lines denote the maximal intervals that may be derived by Prob-EC using a 0.7 threshold, i.e. the set of all consecutive time-points with LTA probability above 0.7. The red line denotes the credible maximal interval computed by PIEC using the same threshold value, and the blue line expresses the ground truth.

Figure 2 illustrates, with the use of a benchmark activity recognition dataset[1], the conditions in which the proposed approach is beneficial. This figure shows a

[1] `http://homepages.inf.ed.ac.uk/rbf/CAVIARDATA1/`

case of probability fluctuation—a steady increase in probability is followed by a noisy observation that reduces dramatically the LTA probability. Subsequently, the probability increases again. PIEC is able to compute a single maximal interval, mitigating the effects of the noisy observation reducing temporarily the LTA probability. In contrast, Prob-EC is directly affected by the noisy observation, creating a series of false negatives between its two maximal intervals. To approximate the interval of PIEC, we would have to lower significantly the threshold value for Prob-EC, creating numerous false positives in other cases.

4 Related Work

The input data of an activity recognition system exhibit various types of uncertainty [2]. One such type is that of incomplete or missing evidence [5]. Additionally, the input events typically have a noise component added to them. Consequently, events are often accompanied by a probability value. Several factors contribute to the corruption of the input events, such as the limited accuracy of sensors and distortion along a communication channel.

A recent survey [2] identified the following classes of methods for handling uncertainty in activity recognition: automata-based methods, probabilistic graphical models, probabilistic/stochastic Petri Nets and approaches based on stochastic (context-free) grammars. The closest line of work to our approach concerns the use of probabilistic graphical models, such as Markov Networks. When used for activity recognition, Markov Networks are combined with first-order logic, in which case they are called Markov Logic Networks (MLN) [9]. The work of Skarlatidis et al. [13] is one of the first attempts to provide a *general* probabilistic framework for activity recognition via MLN. In order to establish such a framework, Skarlatidis and colleagues employed the Event Calculus [8]. They aimed to tackle LTA definition uncertainty, i.e. model imperfect rules expressing LTA. Instead, we built upon a probabilistic Event Calculus handling data uncertainty. Although probabilistic STA can be incorporated into graphical models, correctly encoding their dependencies can be far from obvious, especially with MLN [2].

There are also logic-based approaches to activity recognition that do not (directly) employ graphical models, such as [4, 10, 1, 11]. A key difference between our work and these methods lies in the use of the Event Calculus, which allows us to develop an expressive activity recognition framework, specifying succinctly complex LTA by taking advantage of the built-in representation of inertia. A recently proposed probabilistic Event Calculus is presented in [6]. Our work is complementary to the Event Calculus of [6], as well as that of Skarlatidis et al. [13]. The computation of probabilistic maximal intervals may operate on top of any dialect for point-based probability calculation.

References

1. Massimiliano Albanese, Rama Chellappa, Naresh Cuntoor, Vincenzo Moscato, Antonio Picariello, V. S. Subrahmanian, and Octavian Udrea. PADS: A Probabilistic

Activity Detection Framework for Video Data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2246–2261, 2010.
2. Elias Alevizos, Anastasios Skarlatidis, Alexander Artikis, and Georgios Paliouras. Probabilistic complex event recognition: A survey. *ACM Comput. Surv.*, 50(5):71:1–71:31, 2017.
3. Alexander Artikis, Marek J. Sergot, and Georgios Paliouras. An event calculus for event recognition. *IEEE Trans. Knowl. Data Eng.*, 27(4):895–908, 2015.
4. William Brendel, Alan Fern, and Sinisa Todorovic. Probabilistic event logic for interval-based event recognition. In *The 24th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011, Colorado Springs, CO, USA, 20-25 June 2011*, pages 3329–3336, 2011.
5. Nimrod Busany, Avigdor Gal, Arik Senderovich, and Matthias Weidlich. Interval-based queries over multiple streams with missing timestamps. 2017.
6. Fabio Aurelio D'Asaro, Antonis Bikakis, Luke Dickens, and Rob Miller. Foundations for a probabilistic event calculus. In *Logic Programming and Nonmonotonic Reasoning - 14th International Conference, LPNMR 2017, Espoo, Finland, July 3-6, 2017, Proceedings*, pages 57–63, 2017.
7. A. Dries, A. Kimmig, W. Meert, J. Renkens, G. Van den Broeck, J. Vlasselaer, and L. De Raedt. Problog2: Probabilistic logic programming. In *ECML*, 2015.
8. Robert A. Kowalski and Marek J. Sergot. A logic-based calculus of events. *New Generation Comput.*, 4(1):67–95, 1986.
9. M. Richardson and P. Domingos. Markov logic networks. *Machine Learning*, 62(1–2):107–136, 2006.
10. Joseph Selman, Mohamed R. Amer, Alan Fern, and Sinisa Todorovic. PEL-CNF: Probabilistic event logic conjunctive normal form for video interpretation. In *ICCVW*, pages 680–687. IEEE, 2011.
11. Vinay D. Shet, Jan Neumann, Visvanathan Ramesh, and Larry S. Davis. Bilattice-based Logical Reasoning for Human Detection. In *(CVPR)*, pages 1–8. IEEE Computer Society, 2007.
12. Anastasios Skarlatidis, Alexander Artikis, Jason Filipou, and Georgios Paliouras. A probabilistic logic programming event calculus. *TPLP*, 15(2):213–245, 2015.
13. Anastasios Skarlatidis, Georgios Paliouras, Alexander Artikis, and George A. Vouros. Probabilistic event calculus for event recognition. *ACM Trans. Comput. Logic*, 16(2), 2015.