

A Logic Programming Approach to Activity Recognition

Alexander Artikis¹, Marek Sergot² and Georgios Paliouras¹

¹Institute of Informatics & Telecommunications, NCSR Demokritos, Athens 15310, Greece

²Department of Computing, Imperial College London, SW7 2BT, UK
{a.artikis, paliourg}@iit.demokritos.gr, mjs@doc.ic.ac.uk

ABSTRACT

We have been developing 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 detected on video frames. The output of our system is a set of recognised long-term activities, which are pre-defined temporal combinations of short-term activities. The constraints on the short-term activities that, if satisfied, lead to the recognition of a long-term activity, are expressed using a dialect of the Event Calculus. We illustrate the expressiveness of the dialect by showing the representation of several typical complex activities. Furthermore, we present a detailed evaluation of the system through experimentation on a benchmark dataset of surveillance videos.

Categories and Subject Descriptors

I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods

General Terms

Experimentation

Keywords

event recognition, event calculus

1. INTRODUCTION

A common approach to human activity recognition separates low-level from high-level 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’. The output of the latter type of recognition is a set of ‘long-term activities’, ie pre-defined temporal combinations of short-term activities. We focus on high-level recognition.

We define a set of long-term activities of interest, such as ‘fighting’ and ‘meeting’, as temporal combinations of short-term activities — eg, ‘walking’, ‘running’, and ‘inactive’ (standing still) — using a logic programming implementation of the Event Calculus [9]. More precisely, we employ the Event Calculus to express the temporal constraints on a set of short-term activities that, if satisfied, lead to the recognition of a long-term activity.

We presented preliminary results on activity recognition from video content in [2]. (In [3] we described some initial steps towards automatically constructing activity definitions using machine learning techniques — the use of such techniques is out of the scope of this paper.) In this paper we extend our previous work in the following ways. First, we use a more efficient Event Calculus dialect and implementation to compute the intervals of long-term activities. Second, we illustrate the expressiveness of the proposed Event Calculus dialect by presenting several complex activity definitions. We are able to construct much more succinct representations of activity definitions for video surveillance than we had in our earlier work. Third, we present a more detailed and informative evaluation of the Event Calculus on activity recognition. We show through experimentation how incomplete short-term activity narratives, inconsistent annotation of short-term and long-term activities, and a limited dictionary of short-term activities and context variables affect recognition accuracy. Fourth, we evaluate our approach on a dataset with a refined dictionary of short-term activities, in order to validate experimentally our intuition that a finer classification of short-term activities increases, under certain circumstances, the accuracy of long-term activity recognition. Indeed, the refined dictionary of short-term activities — which can be provided by state-of-the-art short-term activity recognition systems — together with the updated long-term activity definitions presented in this paper, lead to much higher Precision and Recall rates.

The remainder of the paper is organised as follows. First, we present the Event Calculus dialect that we employ to formalise activity definitions. Second, we describe the dataset of short-term activities on which we perform long-term activity recognition. Third, we present our knowledge base of long-term activity definitions. Fourth, we present our experimental results. Finally, we discuss related work and outline directions for further research.

2. THE EVENT CALCULUS

Our long-term activity recognition (LTAR) system consists of a logic programming (Prolog) implementation of an Event Calculus dialect. The Event Calculus, introduced by Kowalski and Sergot [9], is a many-sorted, first-order predicate calculus for representing and reasoning about events and their effects. For the dialect used here, hereafter LTAR-EC (event calculus for long-term activity recognition), the time model is linear and it may include real numbers or 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, $F = V$

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EiMM'10, October 25, 2010, Firenze, Italy.

Copyright 2010 ACM 978-1-4503-0176-3/10/10 ...\$10.00.

Table 1: Main Predicates of the LTAR-EC.

Predicate	Meaning
$\text{happensAt}(E, T)$	Event E is occurring at time T
$\text{happensFor}(E, I)$	I is the list of maximal intervals during which event E takes place
$\text{initially}(F = V)$	The value of fluent F is V at time 0
$\text{holdsAt}(F = V, T)$	The value of fluent F is V at T
$\text{holdsFor}(F = V, I)$	I is the list of maximal intervals for which $F = V$ holds continuously
$\text{initiatedAt}(F = V, T)$	At time T a period of time for which $F = V$ is initiated
$\text{terminatedAt}(F = V, T)$	At time T a period of time for which $F = V$ is terminated

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.

An *event description* in LTAR-EC includes axioms that define, among other things, the event occurrences (with the use of the `happensAt` and `happensFor` predicates), 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`, `holdsAt` and `holdsFor` predicates). Table 1 summarises the main predicates of LTAR-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 axioms for `holdsAt` and `holdsFor` are such that, for any fluent F , $\text{holdsAt}(F = V, T)$ if and only if time-point T belongs to one of the maximal intervals of I such that $\text{holdsFor}(F = V, I)$. However, for efficiency the implementation employs different procedures for these two tasks, and various indexing techniques to reduce search and improve efficiency further. Briefly, to compute $\text{holdsFor}(F = V, I)$, we find all time-points T_i in which $F = V$ is initiated, and then, for each T_i , we compute the first time-point after T_i in which $F = V$ is terminated. If the list of initiating time-points is generated in sorted order, which is easy to arrange, both steps can make effective use of indexing. In particular, if the list of initiating time-points contains an adjacent pair $\dots, T_i, T_{i+1}, \dots$ then the terminating time-point corresponding to T_i must occur between T_i and T_{i+1} . In outline, the indexing works as follows.

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

$$\text{holdsAt}(F = V, T) \leftarrow \text{initiatedAt}(F = V, T_s), \text{ not broken}(F = V, T_s, T) \quad (1)$$

$$\text{broken}(F = V, T_s, T) \leftarrow \text{terminatedAt}(F = V, T_f), T_s < T_f < T \quad (2)$$

$$\text{broken}(F = V_1, T_s, T) \leftarrow \text{initiatedAt}(F = V_2, T_f), V_1 \neq V_2, T_s < T_f < T \quad (3)$$

not in rule (1) represents ‘negation by failure’, which provides a form of default persistence (‘inertia’) of fluents.

According to rule (2), a period of time for which $F = V$ holds is *broken* at T_f if $F = V$ is terminated at time T_f . According to rule (3), if $F = V_2$ is initiated at T_f then effectively $F = V_1$ is terminated at time T_f , for all other possible values V_1 of F . Rule (3) ensures therefore that a fluent cannot have more than one value at any time.

Besides the general, domain-independent rule

$\text{initiatedAt}(F = V, 0) \leftarrow \text{initially}(F = V)$, the definitions of `initiatedAt` and `terminatedAt` are domain specific. One common form of rule for `initiatedAt`, eg, has the general form:

$$\text{initiatedAt}(F = V, T) \leftarrow \text{happensAt}(Ev, T), \text{ Conditions}[T] \quad (4)$$

where $\text{Conditions}[T]$ is some set of further conditions referring to time T . Concrete examples of `initiatedAt` rules are provided in the section that follows.

To explain what we mean by indexing, note that clauses (1), (2) and (3) can be written equivalently as follows:

$$\text{holdsAt}(F = V, T) \leftarrow \text{initiatedAt}(F = V, 0, T_s, T), \text{ not broken}(F = V, T_s, T) \quad (5)$$

$$\text{broken}(F = V, T_{min}, T_{max}) \leftarrow \text{terminatedAt}(F = V, T_{min}, T_f, T_{max}) \quad (6)$$

$$\text{broken}(F = V, T_{min}, T_{max}) \leftarrow \text{initiatedAt}(F = V_2, T_{min}, T_f, T_{max}), V_1 \neq V_2 \quad (7)$$

when every rule of the form (4) is transformed into the form:

$$\text{initiatedAt}(F = V, T_{min}, T, T_{max}) \leftarrow \text{happensAt}(Ev, T_{min}, T, T_{max}), \text{ Conditions}[T] \quad (8)$$

The extra arguments in `initiatedAt`, `terminatedAt`, and `happensAt` specify the range of time-points T_{min} and T_{max} between which the time-point T of interest must occur. Thus $\text{happensAt}(Ev, T_{min}, T, T_{max})$ iff $\text{happensAt}(Ev, T)$ and $T_{min} < T < T_{max}$.

Our implementation automatically transforms `initiatedAt` rules of the form (4) into the form (8) on compilation, in a process transparent to the user. The advantage is that in Prolog execution of `holdsAt` and `holdsFor`, when $\text{happensAt}(Ev, T_{min}, T, T_{max})$ is called, Ev is always ground (variable-free), which exploits Prolog’s built-in indexing when searching for occurrences of event Ev . But much more importantly, T_{min} and T_{max} are also always ground which means that the storage of `happensAt` data can be indexed to exploit this and reduce search very significantly.

`terminatedAt` and other forms of `initiatedAt` rules are handled similarly. We omit these and other details for lack of space. The complete code for LTAR-EC is available upon request.

3. SHORT-TERM ACTIVITIES

Our long-term activity recognition system (LTAR) includes long-term activity definitions in LTAR-EC. The input to LTAR is a symbolic representation of short-term activities. The output of LTAR is a set of recognised long-term activities. In [2, 3] we used the first dataset of the CAVIAR project¹ to perform long-term activity 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 in order to provide the ground truth for both short-term and long-term activities.

Our preliminary experiments with this dataset, however, showed that the limited dictionary of short-term activity types compromised the recognition of some long-term activities — it was often impossible to distinguish between certain long-term activities. To overcome this problem, in the context of this paper we introduced in the CAVIAR dataset a short-term activity for ‘abrupt motion’: we manually edited the annotation of the CAVIAR videos by changing, when necessary, the

¹<http://groups.inf.ed.ac.uk/vision/CAVIAR/CAVIARDATA1/>

label of a short-term activity to ‘abrupt motion’. This is a form of short-term activity that is recognised by some state-of-the-art recognition systems, such as [8]. A person is said to exhibit an ‘abrupt motion’ activity if he moves abruptly and his position in the global coordinate system does not change significantly — if it did then the short-term activity would be classified as ‘running’. For this set of experiments, therefore, the input to LTAR is:

- (i) The short-term activities abrupt motion, walking, running, active (non-abrupt body movement in the same position) and inactive (standing still), together with their time-stamps, ie the video frame in which that short-term activity took place. These activities are mutually exclusive. This type of input is represented by means of the `happensAt` predicate — eg, `happensAt(abrupt(id6), 15560)` expresses that *id₆* moved abruptly at video frame (time-point) 15560. Short-term activities are represented as events in the Event Calculus in order to use the `initiatedAt` and `terminatedAt` predicates for expressing the conditions in which these activities initiate and terminate a long-term activity.
- (ii) The coordinates of the tracked people and objects as pixel positions at each time-point. The coordinates are represented with the use of the `holdsAt` predicate — eg, `holdsAt(coord(id2)=(14, 55), 10600)` expresses that the coordinates of *id₂* are (14, 55) at time-point (frame number) 10600.
- (iii) The first and the last time a person or object is tracked (‘appears’/‘disappears’). This type of input is represented using the `happensAt` predicate. Eg, `happensAt(appear(id10), 300)` expresses that *id₁₀* is first tracked at time-point (frame number) 300.

Given this input, LTAR recognises the following long-term activities: a person leaving an object, a person being immobile, people meeting, moving together, or fighting. Long-term activities are represented as Event Calculus fluents in order to use the `holdsFor` predicate for computing the intervals of these activities. Eg, `holdsFor(moving(id1, id3) = true, [(0, 40), (340, 380)])` states that *id₁* was moving together with *id₃* in the intervals (0, 40), and (340, 380).

To recognise long-term activities, LTAR processes the input information as follows. First, given input type (i), ie the short-term activities detected at each time-point and recorded using the `happensAt` predicate, LTAR computes the maximal duration of each short-term activity, and represents it using the `happensFor` predicate. Eg, `happensFor(walking(id5), [(40, 400), (600, 720)])` expresses that the maximal intervals for which *id₅* was walking are (40, 400) and (600, 720). `appear(A)` and `disappear(A)` are instantaneous events. (They occur at one time point.) Second, given input type (ii), LTAR computes the distance between two tracked entities and compares the distance with pre-defined thresholds. Eg, `holdsAt(close(id3, id5, 30) = true, 80)` expresses that *id₃* is ‘close’ to *id₅* at time 80 in the sense that their distance is at most 30 pixel positions. Further, LTAR computes the maximal intervals for which two tracked entities are ‘close’ — eg, `holdsFor(close(id3, id5, 24) = true, [(40, 80)])` states that (40, 80) is the maximal interval for which the distance between *id₃* and *id₅* is continuously at most 24 pixel positions.

Long-term activity recognition is based on a knowledge base of long-term activity definitions. Next we present example definition fragments of LTAR’s knowledge base.

4. LONG-TERM ACTIVITY DEFINITIONS

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

$$\begin{aligned} & \text{initiatedAt}(\text{leaving_object}(P, Obj) = \text{true}, T) \leftarrow \\ & \quad \text{happensAt}(\text{appear}(Obj), T), \\ & \quad \text{happensAt}(\text{inactive}(Obj), T), \\ & \quad \text{holdsAt}(\text{close}(P, Obj, 30) = \text{true}, T), \\ & \quad \text{holdsAt}(\text{person}(P) = \text{true}, T), \\ & \quad \text{happensAt}(\text{appear}(P), T_0), T_0 < T \end{aligned} \quad (9)$$

$$\begin{aligned} & \text{terminates}(\text{leaving_object}(P, Obj) = \text{true}, T) \leftarrow \\ & \quad \text{happensAt}(\text{disappear}(Obj), T) \end{aligned} \quad (10)$$

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, ie ‘appear’, if and only if the person leaves it somewhere. Moreover, objects (as opposed to persons) can exhibit only inactive short-term activity. Accordingly, axiom (9) 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 short-term activity at *T* is ‘inactive’, there is a person *P* ‘close’ to *Obj* at *T*, and *P* has ‘appeared’ at some time earlier than *T*. Recall that `appear(A)` is an event that takes place at the first time *A* is tracked and that the `close(A, B, D)` fluent is true when the distance between *A* and *B* is at most *D* pixel positions. The value of 30 pixel positions was determined from an empirical analysis of CAVIAR.

In CAVIAR there is no explicit information that a tracked entity is a person or an inanimate object. Therefore, in the activity definitions we try to deduce whether a tracked entity is a person or an object given, among others, the detected short-term activities. We defined the fluent `person(P)` to have value true if *P* has exhibited an active, walking, running or abrupt motion short-term activity since *P* ‘appeared’. The value of `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.

Unlike the specification of `person`, it is not clear from the CAVIAR data whether a tracked entity is an object, and for this reason we do not have a fluent explicitly representing that an entity is an object. `person(P) = false` does not necessarily imply that *P* is an object; it may be that *P* is not tracked, or that *P* is an inactive person. Note finally that axiom (9) 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 axiom (9). (This assumption is adequate for CAVIAR. Removing it raises further issues we do not have space to discuss fully here.)

The lack of explicit information that a tracked entity is an inanimate object may compromise recognition accuracy in certain conditions. A discussion about the effects of the limitations of CAVIAR’s dictionary on recognition accuracy will be presented in the next section.

Axiom (10) expresses the conditions in which a `leaving_object` activity ceases to be recognised. In brief, `leaving_object` is terminated when the object in question is picked up. An object that is picked up by someone is no longer tracked — it ‘disappears’ — terminating `leaving_object`.

The long-term activity `immobile` was defined in order to signify that a person is resting in a chair or on the floor,

or has fallen on the floor (eg, fainted). Note that there is no short-term activity in the CAVIAR annotation for the motion of leaning towards the floor or a chair. The absence of such a short-term activity substantially complicates the definition of *immobile*, and, as discussed in the next section, sometimes reduces the accuracy of recognising *immobile*. Below is one of the axioms of the *immobile* definition:

$$\begin{aligned}
&\text{initiatedAt}(\text{immobile}(P) = \text{true}, T) \leftarrow \\
&\quad \text{happensFor}(\text{inactive}(P), \text{Intervals}), \\
&\quad (T, T_1) \in \text{Intervals}, T_1 > T + 54, \\
&\quad \text{holdsAt}(\text{person}(P) = \text{true}, T), \\
&\quad \text{findall}(S, \text{shop}(S), \text{Shops}), \\
&\quad \text{holdsAt}(\text{farS}(P, \text{Shops}, 24) = \text{true}, T)
\end{aligned} \tag{11}$$

immobile(P) is recognised if the following conditions are satisfied. First, P stays inactive for more than 54 frames (see lines 2–3 of axiom (11)). We chose this number of frames, like all other numerical constraints of the definitions, based on empirical analysis of the CAVIAR dataset. Second, P is a person (see line 4 of axiom (11)). With the use of this constraint we distinguish between an inanimate object, which is inactive since it is first tracked, from an immobile person. Third, P is not ‘close’ to a shop (see lines 5–6 of axiom (11)). If P were ‘close’ to a shop then he would have to stay inactive much longer than 54 frames before *immobile* could be recognised. (Those conditions are specified in other axioms defining *immobile* not shown here.) In this way we avoid classifying the activity of browsing a shop as *immobile*. $\text{farS}(A, \text{List}, D)$ is true when A is more than D pixel positions away from every element of the *List*.

immobile(P) is terminated when P starts walking, running or ‘disappears’, ie he is no longer tracked by the video cameras. The relevant axioms for *terminatedAt* are straightforward and are not shown here.

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 neither running nor moves abruptly, and the distance between them is at most 25 pixel positions. In the CAVIAR annotations, this interaction phase can be seen as some form of greeting (eg, a handshake). The rule below shows one set of conditions in which *meeting* is initiated:

$$\begin{aligned}
&\text{initiatedAt}(\text{meeting}(P_1, P_2) = \text{true}, T) \leftarrow \\
&\quad \text{holdsAt}(\text{close}(P_1, P_2, 25) = \text{true}, T), \\
&\quad \text{holdsAt}(\text{person}(P_1), T), \\
&\quad \text{happensAt}(\text{inactive}(P_1), T), \\
&\quad \text{holdsAt}(\text{person}(P_2), T), \\
&\quad \text{not happensAt}(\text{running}(P_2), T), \\
&\quad \text{not happensAt}(\text{abrupt}(P_2), T)
\end{aligned} \tag{12}$$

meeting is terminated when the two people walk away from each other, or one of them starts running, moves abruptly, or ‘disappears’. The formalisation is straightforward and so omitted here.

The activity *moving* was defined in order to recognise whether two people are walking along together. This activity, like the activities presented so far, could be formalised in terms of *initiatedAt*/*terminatedAt* predicates to specify the conditions in which *moving* starts/ceases to be recognised, and then using the domain-independent axioms of *holdsFor* to compute the maximal intervals of this activity: *moving* is initiated when two people are walking and are ‘close’ to each other, and terminated when the people walk away from each other, when they stop moving, ie become active or inactive, when one of them starts running, moves abruptly, or ‘disappears’.

A considerably more concise representation of *moving*, however, can be given directly in terms of *holdsFor*:

$$\begin{aligned}
&\text{holdsFor}(\text{moving}(P_1, P_2) = \text{true}, \text{MovingI}) \leftarrow \\
&\quad \text{holdsFor}(\text{close}(P_1, P_2, 34) = \text{true}, \text{CloseI}), \\
&\quad \text{happensFor}(\text{walking}(P_1), \text{WalkingI}_1), \\
&\quad \text{happensFor}(\text{walking}(P_2), \text{WalkingI}_2), \\
&\quad \text{intersect_all}([\text{WalkingI}_1, \text{WalkingI}_2, \text{CloseI}], \text{MovingI})
\end{aligned} \tag{13}$$

CloseI are the maximal intervals in which the distance between P_1 and P_2 is continuously at the recorded 34 pixel positions. We compute these intervals using the recorded trajectories of P_1 and P_2 given as input to LTAR. *intersect.all* computes the intersection of a list of intervals. The implementation of *intersect.all* and other constructs manipulating intervals is available with the source code of LTAR-EC. According to axiom (13), the maximal intervals in which P_1 and P_2 are *moving* together are produced by the intersection of the intervals in which P_1 is ‘close’ to P_2 , P_1 is walking and P_2 is walking.

As in the case of *moving*, we could also have formalised *leaving_object*, *immobile* and *meeting* directly in terms of *holdsFor* (as opposed to representing these activities in terms of *initiatedAt* and *terminatedAt* and then using the domain-independent axioms of *holdsFor* to compute their maximal intervals). However, formalising *leaving_object*, *immobile* and *meeting* directly in terms of *holdsFor* is not more concise than formalising these activities in terms of *initiatedAt* and *terminatedAt*. For *leaving_object*, *immobile* and *meeting* it is much simpler to identify the conditions in which these activities are initiated and terminated, than identifying all possible conditions in which these activities hold.

The last definition of LTAR’s knowledge base concerns the *fighting* activity:

$$\begin{aligned}
&\text{holdsFor}(\text{fighting}(P_1, P_2) = \text{true}, \text{FightingI}) \leftarrow \\
&\quad \text{happensFor}(\text{abrupt}(P_1), \text{AbruptI}), \\
&\quad \text{holdsFor}(\text{close}(P_1, P_2, 24) = \text{true}, \text{CloseI}), \\
&\quad \text{intersect_all}([\text{AbruptI}, \text{CloseI}], \text{AbruptCloseI}) \\
&\quad \text{happensFor}(\text{inactive}(P_2), \text{InactiveI}), \\
&\quad \text{complement}(\text{AbruptCloseI}, \text{InactiveI}, \text{FightingI})
\end{aligned} \tag{14}$$

complement is an implementation of the complement operation. Two people are assumed to be *fighting* if at least one of them is moving abruptly, the other is not inactive, and the distance between them is at most 24 pixel positions. As in the case of *moving*, we expressed the definition of *fighting* directly in terms of *holdsFor* because expressing the conditions in which two people are fighting leads to a more succinct representation than expressing the conditions in which *fighting* is initiated and terminated.

5. EXPERIMENTAL RESULTS

We present experimental results on 28 surveillance videos of the CAVIAR project. These videos contain 26419 frames that were manually annotated by the CAVIAR team in order to provide the ground truth for short-term and long-term activities. We edited the original CAVIAR annotation by introducing a short-term activity for abrupt motion. Table 2 shows the performance of LTAR; it shows, for each long-term activity, the number of True Positives (TP), False Positives (FP) and False Negatives (FN), and the corresponding Recall and Precision values. Long-term activities are recognised with the use of the *holdsFor* Event Calculus predicate.

LTAR achieved high Recall and Precision rates, indicating that it may adequately represent complex activities. Perfect

Table 2: Experimental Results.

Behaviour	TP	FP	FN	Recall	Precision
leaving object	4	0	1	0.8	1
immobile	9	8	0	1	0.52
meeting	6	1	3	0.66	0.85
moving	15	3	2	0.88	0.83
fighting	6	0	0	1	1

Recall and Precision rates were not achieved due to various reasons. One of these reasons concerns the fact that the narrative of short-term activities (produced by manual annotation, in the present experiments) is incomplete. Eg, the single FN concerning *leaving_object* is due to the fact that in the video in question the object was left behind a chair and was not tracked. In other words, the left object never ‘appeared’, it never exhibited a short-term activity.

Another reason for having FP and FN is the lack of consistency in the annotation of the videos; eg, the long-term activity of people walking in the same direction while being ‘close’ to each other is not always classified as *moving* (this type of inconsistency leads to FP concerning the recognition of *moving*), the short-term activity of people being active is sometimes classified as walking (eg, leading to FN in the recognition of *meeting*), and so on.

The most important reason for not achieving perfect Recall and Precision in the CAVIAR dataset concerns the limited dictionary of short-term activities and context variables with which the tracked activity is represented. The recognition of *immobile*, for instance, would be much more accurate if there were a short-term activity for the motion of leaning towards the floor or a chair. In the absence of such an activity, the recognition of *immobile* is primarily based on how long a person is inactive. In the CAVIAR videos a person who falls on the floor or rests in a chair stays inactive for at least 54 frames. Consequently LTAR recognises *immobile* if, among other things, a person stays inactive for at least 54 frames. There are situations, however, in which a person stays inactive for more than 54 frames and has not fallen on the floor or sat in a chair: people watching a fight, or just staying inactive waiting for someone. It is in those situations that we have the FP concerning *immobile*.

For similar reasons we did not achieve perfect Recall and Precision in the recognition of *meeting*; it is impossible to define this activity precisely due to the absence of a short-term activity for ‘greeting’.

A particular refinement of CAVIAR’s dictionary — the introduction of a short-term activity for abrupt motion — considerably increased LTAR’s recognition accuracy. More precisely, compared to our earlier results [2, 3], the introduction of abrupt motion reduced the number of FP regarding *moving* and *meeting*. In the original annotation of CAVIAR, the short-term activities of people *fighting* were sometimes classified as walking or active. In the first case LTAR incorrectly recognised *moving*, because two people were walking while being ‘close’ to each other, while in the second case LTAR incorrectly recognised *meeting* (in addition to recognising *fighting*), because two people were active while being ‘close’ to each other. Labelling the short-term activities of people *fighting* as abrupt motion resolved this issue, because abrupt motion does not initiate *moving* or *meeting*.

In addition to increasing the recognition accuracy of *moving* and *meeting*, the introduction of abrupt motion eliminated

FP and FN regarding *fighting*. Moreover, the introduction of abrupt motion did not increase FP or FN in the recognition of the other long-term activities.

Similar to introducing abrupt motion, we could have enhanced CAVIAR’s dictionary by including activities for greeting a person, falling on the floor, etc, and variables explicitly representing that a tracked entity is an object. We did not do this because we are not aware of any short-term activity recognition systems that detect such activities and explicitly represent the aforementioned type of information. In contrast, there are systems that detect abrupt motion — eg, see [8]. We expect that a finer classification of short-term activities and the addition of context variables such as the one mentioned above, will, under certain circumstances, increase the overall activity recognition accuracy, provided that the long-term activity definitions are updated accordingly.

We should like to point out that the issues identified above do not always compromise recognition accuracy. Eg, the lack of explicit information that a tracked entity is an object did not affect the recognition accuracy of *leaving_object* in the 28 CAVIAR videos. This lack of information would have led to FP in the recognition of *leaving_object* in certain conditions, but these conditions did not arise in the CAVIAR videos. Similarly, the lack of consistency in the annotation of activities, and the incompleteness of short-term activity narratives do not always lead to FP or FN. In any case, in the next section we discuss ways of addressing these issues.

Concerning recognition efficiency, we were able to recognise each long-term activity in less than 1 second CPU time, given as input around 1800 temporally sorted short-term activities representing, on average, a CAVIAR video, on an Intel Core i7 920@2.67GHz with 6 GB RAM running Linux Kernel 2.6. Ways to further improve recognition efficiency are presented in the following section.

6. DISCUSSION

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

A well-known system for activity recognition is the Chronicle Recognition System (CRS)². A ‘chronicle’ can be seen as a long-term activity — it is expressed in terms of a set of events (short-term activities in our example), linked together by time constraints, and, possibly, a set of context constraints. The language of CRS relies on a reified temporal logic, where propositional terms are related to time-points or other propositional terms. Time is considered as a linearly ordered discrete set of instants. The language includes predicates for persistence and event absence. Details about CRS may be found on the web page of the system and [5].

The CRS language does not allow mathematical operators in the constraints of atemporal variables. Consequently, the computation of the distance between two people/objects cannot be computed. CRS, therefore, cannot be directly used for activity recognition in video surveillance applications. More generally, CRS cannot be directly used for activity recognition in applications requiring any form of spatial reasoning, or any other type of atemporal reasoning. These limitations could be overcome by developing a separate tool for atemporal reasoning that would be used by CRS whenever this form of reasoning was required. To the best of our knowledge, such extensions of CRS are not available. Clearly, the computa-

²<http://crs.elibel.tm.fr/>

tional efficiency of CRS, which is one of the main advantages of using this system for activity recognition, would be compromised by the integration of an atemporal reasoner.

Hakeem and Shah [7] have presented a hierarchical event representation for analysing videos. The temporal relations between the sub-events of an event definition (or activity, in the terminology of this paper) are represented using the interval algebra of Allen and Ferguson [1] and an extended form of the CASE representation [6] originally used for the syntactic analysis of natural languages.

In our approach to activity recognition, the availability of the full power of logic programming is one of the main attractions of employing the Event Calculus as the temporal formalism. It allows activity definitions to include not only complex temporal constraints — LTAR-EC is at least as expressive as the CRS language and the extended CASE representation with respect to temporal representation — but also complex atemporal constraints. Moreover, when necessary more expressive Event Calculus dialects may be adopted (see, eg, [10]).

Shet et al. have presented a logic programming approach to activity recognition. See [12, 13] for two recent publications. These researchers have presented activity definitions concerning theft, entry violation, unattended packages, and so on. A distinguishing feature of our approach with respect to this line of work concerns the fact that we use the Event Calculus for temporal representation and reasoning. The temporal aspects of the definitions of Shet, Davis et al. are crudely represented — eg, there are no rules for computing the intervals in which a long-term activity takes place. In contrast, the Event Calculus has built-in axioms for complex temporal representation, including the formalisation of inertia, durative events, events with delayed effects, etc, which help considerably the system designer develop activity definitions. Shet and colleagues stated that “[i]n the future we would like to extend this system to reason explicitly about temporal information thus helping us [...] to define models for and recognise human activities within a single framework” [13, p.8]. To the best of our knowledge, they have not developed a system for explicit temporal representation and reasoning since.

Shet and colleagues have incorporated in their logic programming framework a mechanism for reasoning over rules and facts that have an uncertainty value attached. We aim to extend our work by allowing for uncertainty values in the rules of activity definitions in order to address, to a certain extent, the issues arising from incomplete short-term activity narratives, inconsistent annotation of short-term and long-term activities, and a limited dictionary of short-term activities and context variables.

Paschke and colleagues [11] have also proposed the use of an Event Calculus dialect for event recognition. This dialect and LTAR-EC have numerous differences. For example, unlike LTAR-EC, there is no support in the dialect of Paschke et al for multi-valued fluents — only Boolean fluents are considered. Moreover, the treatment of intervals is quite different. The Event Calculus dialect of Paschke and colleagues, for instance, does not include axioms for recognising an ‘on-going’ long-term activity, ie a activity that started taking place at some earlier time-point and still holds. There are also very significant differences in the implementations.

Apart from the numerous differences in expressiveness and implementation, a key contribution of the work presented here, as we see it, is that we have illustrated the expressive-

ness of the Event Calculus for complex activity recognition on a benchmark example, showed a range of different types of definition, and evaluated the adequacy of our representation empirically. We expect that the example itself will be a valuable resource in future uses of the Event Calculus for activity/event recognition.

A logic programming approach to activity recognition has, among others, the advantage that machine learning techniques can be directly employed for developing/refining activity definitions. An area of current work is the use of abductive and inductive logic programming techniques for learning activity definitions. Details about this line of work are given in [3].

LTAR-EC does not currently store the outcome of query computation, ie the intervals of the recognised activities. Consequently, LTAR-EC often performs unnecessary computations, re-computing activity intervals that it already computed but did not store. We are currently experimenting to find the most effective options for caching in LTAR-EC, including those presented in [4].

Acknowledgements

This work was supported partly by the EU PRONTO Project (FP7-ICT 231738).

7. REFERENCES

- [1] J. Allen and G. Ferguson. Actions and events in interval temporal logic. *Journal of Logic and Computation*, 4(5):531–579, 1994.
- [2] A. Artikis and G. Paliouras. Behaviour recognition using the event calculus. In *Artificial Intelligence Applications & Innovations*. Springer, 2009.
- [3] A. Artikis, A. Skarlatidis, and G. Paliouras. Behaviour recognition from video content using the event calculus. *International Journal of AI Tools*, 19(2), 2010.
- [4] L. Chittaro and A. Montanari. Efficient temporal reasoning in the cached event calculus. *Computational Intelligence*, 12(3):359–382, 1996.
- [5] C. Dousson and P. Le Maigat. Chronicle recognition improvement using temporal focusing and hierarchisation. In *IJCAI*, pages 324–329, 2007.
- [6] C. Fillmore. The case for CASE. In E. Bach and R. Harms, editors, *Universals in Linguistic Theory*, pages 1–88. Holt, Rinehart, and Winston, 1968.
- [7] A. Hakeem and M. Shah. Learning, detection and representation of multi-agent events in videos. *Artificial Intelligence*, 171(8–9):586–605, 2007.
- [8] D. Kosmopoulos, P. Antonakaki, K. Valasoulis, A. Kesidis, and S. Perantonis. Human behaviour classification using multiple views. In *Hellenic Conference on Artificial Intelligence*, LNAI 5138. Springer, 2008.
- [9] R. Kowalski and M. Sergot. A logic-based calculus of events. *New Generation Computing*, 4(1):67–96, 1986.
- [10] R. Miller and M. Shanahan. The event calculus in a classical logic. *JETA*, 4(16), 2000.
- [11] A. Paschke and M. Bichler. Knowledge representation concepts for automated SLA management. *Decision Support Systems*, 46(1):187–205, 2008.
- [12] V. Shet, D. Harwood, and L. Davis. VidMAP: video monitoring of activity with Prolog. In *Advanced Video and Signal Based Surveillance*. IEEE, 2005.
- [13] V. Shet, J. Neumann, V. Ramesh, and L. Davis. Bilattice-based logical reasoning for human detection. In *Computer Vision and Pattern Recognition*. IEEE, 2007.