

Distributed Online Learning of Event Definitions

Nikos Katzouris¹, Alexander Artikis^{2,1}, and Georgios Paliouras¹

¹ National Center for Scientific Research “Demokritos”, Athens, Greece

² Department of Maritime Studies, University of Piraeus, Piraeus, Greece
{nkatz, a.artikis, paliourg}@iit.demokritos.gr

Abstract. Logic-based event recognition systems infer occurrences of events in time using a set of event definitions in the form of first-order rules. The Event Calculus is a temporal logic that has been used as a basis in event recognition applications, providing among others, direct connections to machine learning, via Inductive Logic Programming (ILP). OLED is a recently proposed ILP system that learns event definitions in the form of Event Calculus theories, in a single pass over a data stream. In this work we present a version of OLED that allows for distributed, online learning. We evaluate our approach on a benchmark activity recognition dataset and show that we can significantly reduce training times, exchanging minimal information between processing nodes.

1 Introduction

Event recognition systems [5] process sequences of *simple events*, such as sensor data, and recognize *complex events*, i.e. events that satisfy some pattern. Logic-based systems for event recognition typically use a knowledge base of first-order rules to represent complex event patterns and a reasoning engine to detect such patterns in the incoming data. The Event Calculus (EC) [16] has been used as the basis for event recognition systems [1], offering direct connections to machine learning, via Inductive Logic Programming (ILP) [4].

Event recognition applications deal with noisy data streams. Methods that learn from such streams typically build a decision model by a single pass over the input [9]. OLED (Online Learning of Event Definitions) [15] is an ILP system that learns event definitions in the form of EC theories in a single pass over a relational data stream. In this work we present an extension of OLED, that allows for learning a theory in an online and parallel fashion, from disjoint, possibly geographically distributed data streams (it therefore assumes no shared memory). Our approach is based on a simple parallelization scheme of the core OLED functionality and it is appropriate for distributed learning. OLED learns clauses in top-down manner, by gradually specializing an over-general clause. Its single-pass strategy is based on the Hoeffding bound [12], a tool that allows to build decision models by estimating their quality on a small subset of the input. OLED uses the Hoeffding bound to estimate the quality of candidate specializations of a clause on subsets of the input stream. In the proposed parallelization strategy, clauses are evaluated independently on disjoint data streams and their scores are combined whenever a specialization decision must be made. We present an evaluation of our approach on an a benchmark activity recognition dataset and show that

Predicate	Predicate Meaning	Axioms
$\text{happensAt}(E, T)$	Event E occurs at time T	$\text{holdsAt}(F, T + 1) \leftarrow$ $\text{initiatedAt}(F, T).$ (1)
$\text{initiatedAt}(F, T)$	At time T a period of time for which fluent F holds is initiated	
$\text{terminatedAt}(F, T)$	At time T a period of time for which fluent F holds is terminated	$\text{holdsAt}(F, T + 1) \leftarrow$ $\text{holdsAt}(F, T),$ (2) $\text{not terminatedAt}(F, T).$
$\text{holdsAt}(F, T)$	Fluent F holds at time T	

Table 1. The basic predicates and domain-independent axioms of the EC dialect.

we can significantly reduce training times, exchanging minimal information between processing nodes. This work, therefore, paves the way for relational learning in high-velocity data streams.

The rest of this paper is structured as follows: In Section 2 we present some background on the EC and provide a running example. In Section 3 we present OLEd and in Section 4 we present its distributed version. In Section 5 we present our experimental results, while in Section 6 we discuss related work. Finally, in Section 7 we discuss some directions for future work and conclude.

2 Background and Running Example

The Event Calculus (EC) [16] is a temporal logic for reasoning about events and their effects. Its ontology consists of *time points* (integer numbers); *fluents*, i.e. properties that have different values in time; and events, i.e. occurrences in time that may alter fluents' values. The axioms of the EC incorporate the *common sense law of inertia*, according to which fluents persist over time, unless they are affected by an event. We use a simplified version of the EC that has been shown to suffice for event recognition [1]. The basic predicates and its domain-independent axioms are presented in Table 1. Axiom (1) states that a fluent F holds at time T if it has been initiated at the previous time point, while Axiom (2) states that F continues to hold unless it is terminated. Definitions for $\text{initiatedAt}/2$ and $\text{terminatedAt}/2$ predicates are given in an application-specific manner by a set of *domain-specific* axioms.

We illustrate our approach using the task of activity recognition, as defined in the CAVIAR project³. The CAVIAR dataset consists of videos where actors perform some activities. Manual annotation (performed by the CAVIAR team) provides ground truth for two activity types. The first type corresponds to simple events and consists of knowledge about the activities of a person at a certain video frame/time point, such as *walking*, or *standing still*. The second activity type corresponds to complex events and consists of activities that involve more than one person, e.g. two people *meeting each other*, or *moving together*. The goal is to recognize complex events as combinations of simple events and additional domain knowledge, such as a person's direction and position.

Table 2(a) presents some example CAVIAR data, consisting of a narrative of simple events in terms of $\text{happensAt}/2$, expressing people's short-term activities, and context properties in terms of $\text{holdsAt}/2$, denoting people's coordinates and direction. Table 2(a)

³ <http://homepages.inf.ed.ac.uk/rbf/CAVIARDATA1/>

(a)	(b)
Narrative for time 1:	Two Domain-specific axioms:
happensAt(walk(id_1), 1)	initiatedAt(moving(X, Y), T) \leftarrow
happensAt(walk(id_2), 1)	happensAt(walk(X), T),
holdsAt(coords(id_1 , 201, 454), 1)	happensAt(walk(Y), T),
holdsAt(coords(id_2 , 230, 440), 1)	distLessThan(X, Y , 25, T),
holdsAt(direction(id_1 , 270), 1)	dirLessThan(X, Y , 45, T).
holdsAt(direction(id_2 , 270), 1)	
Annotation for time 1:	Annotation for time 2:
not holdsAt(move(id_1, id_2), 1)	terminatedAt(move(X, Y), T) \leftarrow
	happensAt(inactive(X), T),
	distMoreThan(X, Y , 30, T).

Table 2. (a) Example data from activity recognition. E.g., at time point 1 person id_1 is walking, her (x, y) coordinates are (201, 454) and her direction is 270° . The annotation for the same time point states that persons id_1 and id_2 are not moving together, in contrast to the annotation for time point 2. **(b)** An example of two domain-specific axioms in the EC. E.g. the first clause dictates that moving of two persons X and Y is initiated at time T if both X and Y are walking at time T , their euclidean distance is less than 25 and their difference in direction is less than 45° .

also shows the annotation of complex events (long-term activities) for each time-point in the narrative. Negated complex events’ annotation is obtained via the closed world assumption (although both positive and negated annotation atoms are presented in Table 2, to avoid confusion). Table 2(b) presents two domain-specific axioms in the EC.

Our goal is to learn definitions of complex events in terms of initiation and termination conditions, as in Table 2(b). In the Learning from Interpretations [3] ILP setting that we use in this work, each training example is an interpretation, i.e. a set of true ground atoms, as in Table 2(a). Given a set of training interpretations \mathcal{I} and some background theory B , which in our case consists of the domain-independent axioms of the EC, the goal is to find a theory H , such that $B \cup H$ covers as many $I \in \mathcal{I}$ as possible, where $B \cup H$ covers I when I is a model of $B \cup H$. Although different semantics are possible, in this work a “model” is a stable model [10].

3 The OLED System

OLED [15] learns a theory by joining together independently-constructed clauses, each of which is learnt in an online fashion. OLED relies on the Hoeffding bound [12] to approximate the quality of a clause on the entire input using only a subset of the input. Given a random variable X with range in $[0, 1]$ and an observed mean \bar{X} of its values after n independent observations, the Hoeffding Bound states that, with probability $1 - \delta$, the true mean \hat{X} of the variable lies in an interval $(\bar{X} - \epsilon, \bar{X} + \epsilon)$, where $\epsilon = \sqrt{\frac{\ln(1/\delta)}{2n}}$. In other words, the true average can be approximated by the observed one with probability $1 - \delta$, given an error margin ϵ that becomes smaller as the number of observations n increases.

OLED learns a clause in a top-down fashion, by specializing it using literals from a bottom clause [4]. The Hoeffding bound is utilized in the specialization process as follows: Given a clause evaluation function G and some clause r , OLED evaluates r and all of its candidate specializations on training examples that stream-in. Assume that

after n training examples from the input stream, r_1 is r 's specialization with the highest observed mean G -score \bar{G} and r_2 is the second-best one, i.e. $\Delta\bar{G} = \bar{G}(r_1) - \bar{G}(r_2) > 0$. Then by the Hoeffding bound we have that for the true mean of the scores' difference $\Delta\hat{G}$ it holds that $\Delta\hat{G} > \Delta\bar{G} - \epsilon$, with probability $1 - \delta$, where $\epsilon = \sqrt{\frac{\ln(1/\delta)}{2n}}$. Hence, if $\Delta\bar{G} > \epsilon$ then $\Delta\hat{G} > 0$, implying that r_1 is indeed the best specialization, with probability $1 - \delta$. In order to decide which specialization to select, it thus suffices to accumulate examples from the input stream until $\Delta\bar{G} > \epsilon$. These examples need not be stored or reprocessed. Each example is processed once to extract the necessary statistics for calculating G -scores and it is subsequently discarded, thus giving rise to an online (single-pass) clause construction strategy. To ensure that no clause r is replaced by a specialization of lower quality, r itself is also considered as a potential candidate along with its specializations, ensuring that specializing r is a better decision, with probability $1 - \delta$, than not specializing it at all.

The default specialization process follows a FOIL-like, hill-climbing strategy, where a single literal is added to a clause at each specialization step. However, OLED supports different specialization strategies as well, e.g. by allowing to simultaneously try all specializations up to a given clause length, or by supporting user-defined, TILDE-like look-ahead specifications [2].

To calculate G -scores, each clause r is equipped with a true positive (TP), a false positive (FP) and a false negative (FN) counter, whose values are updated accordingly as r gets evaluated on training examples that stream-in. True negative counts are not taken into account, since the annotation for complex events is acquired via the closed world assumption. Also, r has an example counter that counts the number of examples on which r has been evaluated so far and is used in the calculation of ϵ in the Hoeffding bound-based search heuristic. Although different scoring functions may be plugged into OLED, in this work we use precision, to score initiation clauses, and recall, to score termination clauses, as in [15]. Moreover, OLED supports a clause pruning mechanism, that allows to remove low-quality clauses (e.g. clauses that have been generated from noisy examples) and a tie-breaking mechanism, that allows to randomly select between equally good specializations. We refer to [15] for more details on these features.

In the general case, a theory learnt by OLED is a collection of clauses constructed with the online mechanism described above. Starting with an empty theory $H = \emptyset$, an initial clause is generated from the first positive example that streams-in, by constructing a bottom clause \perp from that example and adding the empty-bodied clause $r = head(\perp) \leftarrow$ to theory H . From that point on, r is gradually specialized by the addition of literals from \perp to its body. New clauses are added to H whenever existing clauses in H become too specific to account for new incoming examples.

When learning domain-specific axioms in the Event Calculus, the aforementioned generic theory construction strategy must be modified to account for the fact that initiation and termination clauses behave differently w.r.t. encountered TP , FP and FN complex event instances. A description of this behavior is illustrated in Figure 1(A). To handle this behavior, initiation and termination clauses are learnt separately, by two parallel processes, each of which runs the core OLED Algorithm. The input stream is forwarded to each of these processes simultaneously. Figure 1(B) presents the different actions that each learner takes whenever it encounters TP , FP and FN instances.

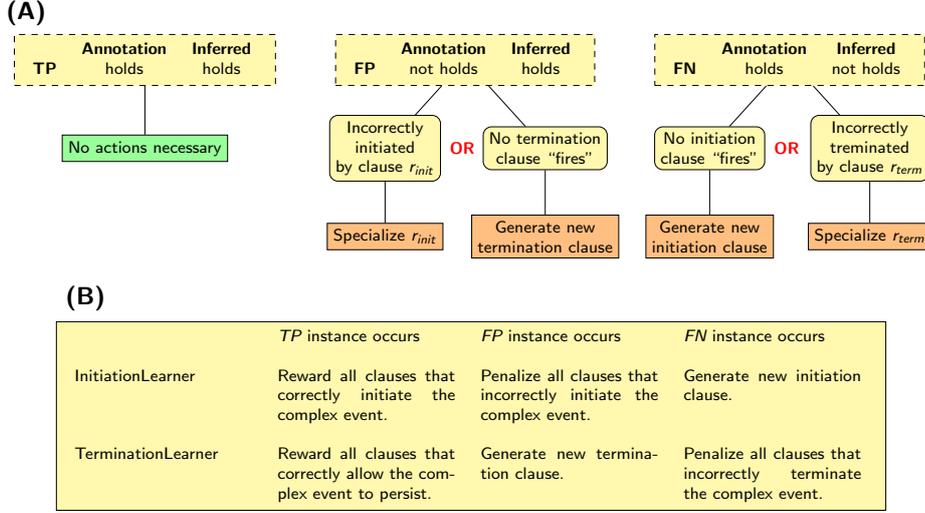


Fig. 1. (A) Different behaviors of initiation and termination clauses w.r.t. to occurrences of *TP*, *FP* and *FN* complex event instances. Dash-lined boxes explain what it means to encounter a *TP*, *FP*, *FN* complex event instance, in terms of (dis)agreement between the actual label of the instance and the one inferred by the theory. Round-cornered boxes describe the causes of *FP*, *FN* occurrences w.r.t. the different types of clause (initiation or termination). Regular boxes at the “leaves” of the tree-like structures indicate proper courses of action in order to eliminate *FP/FN* instances. (B) Actions taken by the two different processes that learn initiation and termination clauses in parallel, w.r.t. *TP*, *FP*, *FN* complex event occurrences. These actions are in accordance with the indicated actions in (A) (leaves of the trees). “Rewarding” a clause refers to increasing the *TP* count of the clause, while “penalizing” a clause refers to increasing its *FP* or *FN* counts. Penalizing clauses reduces their score, it therefore contributes to their specialization after a sufficient number of examples.

4 A Distributed Version of OLED

We now proceed to the description of a distributed version of OLED. Evaluating a clause and its candidate specializations on incoming examples, may be performed in parallel, by distributing the clause evaluation workload across multiple processing nodes that operate on independent data partitions. When needed, e.g. when the Hoeffding test succeeds at some processing node, evaluation results from other nodes may be combined in order to make a more informed decision. We next describe this strategy in more detail.

We assume that learning is performed by a set \mathcal{N} of independent, possibly distributed processing nodes. Each node $N_i \in \mathcal{N}$ handles a separate stream S_i of training examples. The nodes in \mathcal{N} communicate by exchanging messages and they learn a theory H simultaneously, each node using its own training stream. The distributed version of OLED differs from the sequential algorithm in the following respects:

New clause addition: When a node N_i generates a new clause r , it broadcasts r to all other nodes in \mathcal{N} , via a `AddNewClause(r)` message (see Table 3 for the main types of message of distributed OLED). Each node that receives such a message adds clause r

to its own theory and starts scoring r , and its candidate specializations, on its own data. As in the single-core version of OLED, a new clause r consists of an empty-bodied clause $head(\perp_r) \leftarrow$, where \perp_r is a bottom clause generated at N_i .

Clause specialization: When a node N_i is about to specialize a clause r , i.e. when OLED’s Hoeffding test for clause r succeeds, locally at N_i , node N_i sends a `SpecializeRequest(r_{id})` message to all other nodes, where r_{id} is a unique identifier of clause r , common to all copies of r shared among processing nodes. Upon receiving such a message, each node uses r_{id} to retrieve its own evaluation statistics for clause r and its candidate specializations, which are sent over to the requesting node N_i . These statistics consist of TP , FP , FN and E counts for clause r and its candidate specializations, where by E we denote the number of examples on which a clause has been evaluated so far. The received counts for clause r and its specializations are combined with node N_i ’s local counts as follows (we describe the process for clause r only, but it is similar for each one of its specializations). Denoting by TP_r^j , FP_r^j , FN_r^j and E_r^j the respective counts for clause r , received from node $N_j \in \mathcal{N}$, $j \neq i$, the current node N_i updates r ’s counts accordingly, by increasing r ’s local counts with those received from other nodes. For instance, the new TP count for clause r in node N_i becomes $TP_r^i = TP_r^i + \sum_{N_j \in \mathcal{N}} TP_r^j$. FP_r^i , FN_r^i and E_r^i counts are updated in a similar fashion.

Each processing node $N_i \in \mathcal{N}$ maintains a record, for each clause r in its theory and each one of r ’s specializations, that contains the exact counts previously received for them, from each node $N_j \in \mathcal{N}$, $j \neq i$. When node N_i receives a set of new TP_r^j , FP_r^j , FN_r^j and E_r^j counts for clause r from node N_j , $j \neq i$, the respective previous counts are subtracted from the new ones, to avoid over-scoring r with counts that have already been taken into account in previous updates. The same holds for r ’s specializations.

Once individual clause evaluation statistics are combined as described above, node N_i repeats the Hoeffding test for clause r to assess if the test still succeeds after the accumulated counts from all other nodes, for clause r and its specializations, have been taken into account. If it does, clause r is replaced in H , the current theory at node N_i , by its best-scoring specialization r' that results from the Hoeffding test. Then node N_i sends out a `Replace(r_{id} , r')` message to all other nodes, instructing them to also replace r in their own theories with r' . If, on the other hand, the Hoeffding test fails at node N_i after the updated counts are taken into account, clause r is not specialized and all nodes continue evaluating their theories on new incoming examples from their training streams.

Clause pruning: For a clause r to be pruned away, two conditions must hold: First, clause r must be unchanged (not specialized) for a sufficiently long period, which, in the single-core version of OLED, is set to the average number of examples, observed so far in the learning process, for which the Hoeffding test succeeds, i.e. the average value of $n = \mathcal{O}(\frac{1}{\epsilon^2} \ln \frac{1}{\delta})$ that has resulted in clause specializations so far. Second, from that point on where clause r remains unchanged, a sufficiently large number of examples must be seen, in order to use a Hoeffding test to infer that, with probability $1 - \delta$, the quality of clause r is below the pruning threshold.

In the distributed version of OLED, each node uses the above heuristics to decide locally whether a clause r should be pruned. Once it has seen enough data from its

Message	Conditions for message broadcast	Actions upon message receipt
$\text{AddNewClause}(r)$	Generation of clause r .	Add r to local theory.
$\text{SpecializeRequest}(r_{id})$	Clause with id r_{id} is about to be specialized.	Reply to the sender by the local TP, FP, FN, E counts for clause with id r_{id} and for each of its candidate specializations.
$\text{Replace}(r_{id}, r')$	Clause with id r_{id} has been specialized to clause r' .	Replace clause with id r_{id} by r' in local theory.
$\text{PruneRequest}(r_{id})$	Clause with id r_{id} is about to be pruned.	Reply to the sender by the local TP, FP, FN, E counts for clause with id r_{id} , as well as the period for which r remains (locally) unchanged.
$\text{Remove}(r_{id})$	Clause with id r_{id} has been pruned.	Remove clause with id r_{id} from local theory.

Table 3. The main messages exchanged between data processing nodes in distributed OLED.

own stream to make that decision for clause r , it sends a $\text{PruneRequest}(r_{id})$ message to all other nodes. Each node that receives such a message sends back to the requesting node the necessary information (period for which clause r remains unchanged and TP, FP, FN and E_r counts for clause r), which node N_i uses to re-assess whether clause r should be pruned, based on the global view of clause r , obtained by combining r 's separate evaluations from all processing nodes. If node N_i eventually decides to pruned clause r , it sends a $\text{Remove}(r_{id})$ to all other nodes, which instructs them to also remove clause r from their theories.

Algorithm 1 illustrates learning in distributed OLED by displaying the functionality of each processing node.

To ensure that all nodes have the same theory at each point during the learning process, processing nodes often block their execution. For instance, whenever a node N_j sends out a $\text{SpecializeRequest}(r_{id})$, it blocks until it receives the necessary statistics for clause r_{id} from all other nodes. Similarly, whenever a node $N_i, i \neq j$ receives such a message, it replies by sending over to N_j the necessary statistics for clause r_{id} and then blocks its execution, waiting for a “verdict” from N_j . A “verdict” may be a $\text{Replace}(r_{id}, r')$ message, which instructs node N_i to replace clause r_{id} by its specialization r' , or, it may be a Proceed message (this type of message is omitted from Table 3), signifying the fact that the Hoeffding test for clause r_{id} at node N_j failed, after accumulated counts from all other nodes were taken into account, and therefore node N_i may continue processing new examples without altering its current theory. A similar blocking behavior occurs during message passing for clause pruning, to ensure that all all nodes remove (or preserve) a clause in a synchronized fashion.

While exchanging messages during learning, care must be taken to avoid deadlocks and race conditions. An example of such situations is the case where the Hoeffding test

Algorithm 1 OLEDNode($H, G, \mathcal{I}, \mathcal{N}'$)

Input: H : A potentially empty hypothesis; G : A clause evaluation function; \mathcal{I} : A stream of training interpretations; \mathcal{N}' : Set of peer processing nodes.

```

1:  $H := \emptyset$ 
2: while true do
3:   Let  $I$  be the next training interpretation
4:    $G$ -score each clause  $r \in H$  on  $I$ 
5:   if a new clause should be generated then
6:     Generate a new clause  $r$  and add to  $H$ .
7:     Send  $r$  to each node in  $\mathcal{N}'$ .
8:   for each clause  $r \in H$  do
9:     if the Hoeffding test for  $r$  succeeds then
10:      Request the counts for  $r$ , and all of  $r$ 's specializations, from all nodes in  $\mathcal{N}'$ .
11:      Add the received counts to the current ones and repeat the Hoeffding test.
12:      if the Hoeffding test for  $r$  still succeeds then
13:        Replace  $r$  in  $H$  with its best-scoring specialization.
14:        Notify all nodes in  $\mathcal{N}'$  to also replace  $r$  with its best-scoring specialization.
15:      if  $r$  should be pruned then
16:        Request the counts for  $r$  from all nodes in  $\mathcal{N}'$ .
17:        Add the received counts for  $r$  to the current ones.
18:        if  $r$  still should be pruned then
19:          Remove  $r$  from the current theory.
20:          Notify all nodes in  $\mathcal{N}'$  to also remove  $r$  from their theories.
21: return  $H$ 

```

for specializing a single clause r succeeds simultaneously at two different nodes N_i and N_j . In such a case, nodes N_i and N_j send a `SpecializeRequest`(r_{id}) message to each other (in addition to all other nodes), requesting each other's counts for clause r . Subsequently, both nodes enter a "waiting state" as mentioned above, waiting to receive the requested counts in order to proceed with repeating the Hoeffding test. This results in a deadlock, since each node is waiting for the other node's reply. A similar situation may occur with pruning. To avoid such problems, all communication between nodes regarding clause specialization and pruning is mediated by a separate node, whose sole role is to assign priorities in such cases, so that such deadlocks are avoided. When two or more nodes request to specialize or prune a clause (almost) simultaneously, the mediator node prioritizes one of them randomly, enqueueing the other nodes for proceeding at a later time. The enqueued nodes reply to the requests of the prioritized node. Once the prioritized node finishes, the next node from the queue is prioritized. In such cases, a node from the queue abandons its effort to specialize or prune a clause r , if r gets specialized or pruned by the previously prioritized node.

To sum-up, all nodes in \mathcal{N} share a copy of the same theory H at each point in learning. H is learnt in an online fashion, simultaneously from all nodes in \mathcal{N} , each node handling its own training stream. As in the single-core version, initiation and termination clauses are learnt independently, by two separate groups of processes, each of which learns one type of clause and implements the distributed version of OLED described in this section.

		#cores	Time (sec)	Speed-up	F_1 -score	Theory size	#Msgs	Msg size (KB)
(A)	<i>Meeting</i>	1	46	–	0.794	28	–	–
		2	17	2.7	0.802	32	34	97
		4	15	3	0.798	32	52	122
		8	14	3.2	0.800	34	88	182
	<i>Moving</i>	1	73	–	0.731	21	–	–
		2	34	2	0.729	21	41	102
		4	30	2.4	0.731	23	64	188
		8	28	2.6	0.730	23	108	218
(B)	<i>Meeting</i>	1	423	–	0.832	36	–	–
		2	122	3.4	0.832	36	105	247
		4	116	3.6	0.832	36	298	705
		8	110	3.8	0.832	36	518	1527
	<i>Moving</i>	1	514	–	0.758	21	–	–
		2	238	2.1	0.758	21	97	218
		4	225	2.2	0.758	21	265	607
		8	221	2.3	0.758	21	482	1422

Table 4. (A) Experimental results from the CAVIAR dataset; (B) Experimental results from a dataset consisting of 10 copies of CAVIAR.

5 Experimental Evaluation

We present an experimental evaluation of our approach on CAVIAR (described in Section 2), a benchmark dataset for activity recognition. CAVIAR contains 282,067 training interpretations with a mean size of 25 atoms each. OLED is implemented in the Scala programming language. It uses Clingo⁴ as its main reasoning component and Scala’s akka Actors library⁵ to model the behavior of a processing node (Algorithm 1) and implement message passing. The code and data are available online⁶. All experiments were conducted on a Linux machine with a 3.6GHz processor (4 cores and 8 threads) and 16GB of RAM.

The purpose of our experiments was to compare the distributed version of OLED with its monolithic counterpart. We performed learning with 1, 2, 4 and 8 processing cores for acquiring the definitions of two target complex events, related to two persons *meeting each other* or *moving together*. CAVIAR contains 6,272 interpretations in which *moving* occurs and 3,722 in which *meeting* occurs. For the experiments with the distributed version of OLED positive and negative interpretations for the target complex events were evenly distributed across different processing cores. The results were obtained by tenfold cross-validation and are presented in Table 4(A), in the form of averages for training time, F_1 -score and theory size (total number of literals), as well

⁴ <http://potassco.sourceforge.net/>

⁵ <http://akka.io/>

⁶ <https://github.com/nkatzz/OLED>

as average exchanged message number and size. F_1 -scores were obtained by micro-averaging results from each fold.

With respect to the predictive accuracy, multiple-core learning resulted in theories of slightly higher F_1 -score for the *meeting* complex event, as compared to single-core learning. In the monolithic setting, OLED postpones the generation of new clauses, up to the point where existing clauses become too specific to cover new examples. During this time, positive examples which may result in good clauses (recall that OLED learns by “encoding” examples into bottom clauses), are “skipped”, i.e. they are not used for learning new clauses, since they are covered by existing ones. In contrast, the data distribution in the multi-core setting results in cases where interesting examples that would have been missed in the monolithic setting, are actually used for learning. This resulted in OLED learning slightly “richer” theories for the *meeting* complex event in the multi-core setting. A similar effect was not observed for the *moving* complex event, which has a simpler definition than *meeting*.

Regarding training times, OLED achieves a significant speed-up, by moving from monolithic learning to learning with two cores. This is achieved by exchanging very little information between the processing cores (see the last column of Table 4(A)). When additional cores are added, the speed-up is sub-linear, mainly due to the increased number of exchanged messages that result in blocking processing cores.

To test distributed OLED further, we used a larger dataset obtained by sequentially appending to each other 10 copies of CAVIAR, “pushing forward” the time-stamps of the data in each copy. The experimental setting for the $\times 10$ CAVIAR experiment was identical to the one described above for the $\times 1$ CAVIAR experiment. The results are presented in Table 4(B).

Training times in the $\times 10$ CAVIAR experiment follow the same speed-up pattern observed in the $\times 1$ experiment. In this set of experiments, the speed-up factor reaches 3.8, leading to highly reduced training times.

Due to the repetition of the training data in the $\times 10$ CAVIAR experiment, each fold of the tenfold cross-validation process converged to identical theories, regardless of the number of cores used. F_1 -scores are therefore identical for all number of cores, and are also improved as compared to the $\times 1$ CAVIAR experiment. In the latter experiment, good clauses were often constructed “too-late”, from examples that were encountered shortly before the data were exhausted (in which case OLED terminates). Such clauses may be discarded, since OLED uses a “warm-up” period parameter that controls a minimum number of examples on which a clause must be evaluated in order to be included in an output hypothesis. In contrast, in the $\times 10$ CAVIAR experiment such problems were avoided, thanks to the increase in training data size.

6 Related Work

An overview of existing approaches to learning theories in the Event Calculus with ILP may be found at [14, 13] and a discussion on how OLED compares to such approaches may be found at [15, 13]. In this section we mainly discuss distributed/parallel ILP algorithms, for which a substantial amount of work exists in the literature. A thorough review may be found in [8, 22]. Parallel ILP algorithms exploit parallelism across three

main axes [8]: Searching through the hypothesis space in parallel (search parallelism); splitting the training data and learning from data subsets (data parallelism); and evaluating candidate clauses in parallel (evaluation/coverage parallelism).

In [20] the authors present a data-parallel version of a standard set-cover loop: Each processing node learns a fragment of the concept definition from a partition of the data, and then these fragments are exchanged between all nodes. Good-enough clauses are kept by all nodes. A cover removal step is subsequently implemented by each core and the set-cover loop continues. Overall, the approach in [20] learns much faster than a sequential algorithm, achieving super-linear speed-ups. A similar approach is proposed in [7], where the training examples are split across multiple nodes and searched in parallel, while the best rules from each node are “pipe-lined” to all other nodes.

In [22] the authors use a MapReduce-based framework to parallelize the operation of a classical set-cover ILP algorithm towards both evaluation-parallelism and search-parallelism. In the former case, coverage tests of candidate clauses are performed in parallel, on disjoint partitions of the data. In the latter case, bottom clauses (which are generalized to acquire a hypothesis clause) are generated and searched in a concurrent fashion from more than one “seed” examples. The reducer then selects the best hypothesis clause that results from this process. A similar approach for parallel exploration of independent hypotheses has been proposed in [18], while similar approaches towards parallel coverage tests have been proposed in [11, 6]. In [17], the approach of [22] was extended to a framework that is capable of self-regulating the workload of distributing learning costs across multiple nodes.

A main difference of the work presented here from the aforementioned approaches to parallel ILP is that they all rely on iterative ILP algorithms (basically, a set-cover loop), and they all require several passes over the data to compute a hypothesis. In contrast, OLED is an online, single-pass algorithm. In relation to the latter, some work on streaming ILP exists. However, existing approaches are either oriented towards unsupervised tasks like frequent pattern discovery [19], or they rely on propositionalization techniques and off-the-self, online propositional learners [21].

7 Conclusions and Future Work

We presented a distributed version of a recently proposed algorithm for online learning of complex event definitions in the form of domain-specific axioms in the Event Calculus. We also presented an experimental evaluation of our approach on a benchmark dataset for activity recognition, which demonstrates that we can significantly reduce training times. As future work, we aim to evaluate our approach on larger datasets, in terms of in-situ, geographically distributed learning, as in the case of maritime monitoring. Also, we plan to remove the requirement that all processing nodes block their execution, while waiting for replies during message passing.

Acknowledgments. This work was funded by the H2020 project DATACRON.

References

1. Alexander Artikis, Marek Sergot, and Georgios Paliouras. An event calculus for event recognition. *Knowledge and Data Engineering, IEEE Transactions on*, 27(4):895–908, 2015.
2. Hendrik Blockeel and Luc De Raedt. Top-down induction of first-order logical decision trees. *Artificial intelligence*, 101(1):285–297, 1998.
3. Hendrik Blockeel, Luc De Raedt, Nico Jacobs, and Bart Demoen. Scaling up inductive logic programming by learning from interpretations. *Data Mining and Knowledge Discovery*, 3(1):59–93, 1999.
4. Luc De Raedt. *Logical and relational learning*. Springer Science & Business Media, 2008.
5. Opher Etzion and Peter Niblett. *Event processing in action*. Manning Publications Co., 2010.
6. Andreas K Fidjeland, Wayne Luk, and Stephen H Muggleton. Customisable multi-processor acceleration of inductive logic programming. *Latest Advances in Inductive Logic Programming*, pages 123–141, 2014.
7. Nuno A. Fonseca, Fernando M. A. Silva, Vítor Santos Costa, and Rui Camacho. A pipelined data-parallel algorithm for ILP. In *2005 IEEE International Conference on Cluster Computing (CLUSTER 2005), September 26 - 30, 2005, Boston, Massachusetts, USA*, pages 1–10, 2005.
8. Nuno A Fonseca, Ashwin Srinivasan, Fernando Silva, and Rui Camacho. Parallel ilp for distributed-memory architectures. *Machine learning*, 74(3):257–279, 2009.
9. Joao Gama. *Knowledge discovery from data streams*. CRC Press, 2010.
10. Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. Answer set solving in practice. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6(3):1–238, 2012.
11. James H. Graham, C. David Page Jr., and Ahmed H. Kamal. Accelerating the drug design process through parallel inductive logic programming data mining. In *2nd IEEE Computer Society Bioinformatics Conference, CSB 2003, Stanford, CA, USA, August 11-14, 2003*, pages 400–402, 2003.
12. Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
13. Nikos Katzouris. Scalable relational learning for event recognition. *PhD Thesis, University of Athens*, 2017.
14. Nikos Katzouris, Alexander Artikis, and Georgios Paliouras. Incremental learning of event definitions with inductive logic programming. *Machine Learning*, 100(2-3):555–585, 2015.
15. Nikos Katzouris, Alexander Artikis, and Georgios Paliouras. Online learning of event definitions. *TPLP*, 16(5-6):817–833, 2016.
16. Robert Kowalski and Marek Sergot. A logic-based calculus of events. *New Generation Computing*, 4(1):67–95, 1986.
17. Hiroyuki Nishiyama and Hayato Ohwada. Yet another parallel hypothesis search for inverse entailment. In *ILP*, 2015.
18. Hayato Ohwada and Fumio Mizoguchi. Parallel execution for speeding up inductive logic programming systems. In *International Conference on Discovery Science*, pages 277–286. Springer, 1999.
19. Andreia Silva and Cláudia Antunes. Multi-relational pattern mining over data streams. *Data Mining and Knowledge Discovery*, 29(6):1783–1814, 2015.
20. David B Skillicorn and Yu Wang. Parallel and sequential algorithms for data mining using inductive logic. *Knowledge and Information Systems*, 3(4):405–421, 2001.
21. Ashwin Srinivasan and Michael Bain. Relational models with streaming ilp. In *ILP*, 2013.
22. Ashwin Srinivasan, Tanveer A Faruquie, and Sachindra Joshi. Data and task parallelism in ilp using mapreduce. *Machine learning*, 86(1):141–168, 2012.