Scalable Semi-Supervised Structure Learning for Event Recognition

Evangelos Michelioudakis

Doctoral Thesis

Department of Informatics & Telecommunications, National and Kapodistrian University of Athens

Institute of Informatics & Telecommunications, NCSR "Demokritos"

December 2023

Presentation Outline

1 Introduction

- 2 SPLICE: Semi-Supervised Learning for Complex Event Recognition
- 3 SPLICE⁺: Semi-Supervised Learning Combining Structure and Mass-based Predicate Similarity
- 4 Experimental Study
- 5 Conclusions & Future Work



- ► An event represents anything that happens or occurs in time.
- An event may be *instantaneous* or *durative*.
- An event may be *simple* or *complex*.

Event Recognition



SDE: Simple Derived Event CE: Complex Event



Evangelos Michelioudakis



53

Event Recognition INPUT **•** RECOGNITION ► OUTPUT 000 Event Recognition Recognised CEs Streams of SDEs System . . . $leaving_object(id_0, id_1)$ 340 $appear(id_0)$ 340 340 $inactive(id_0)$ 340 $coord(id_0) = (20.88, 11.90)$ CE Definitions 340 $walking(id_2)$ $coord(id_2) = (25.88, 19.80)$ 340340 $active(id_1)$ 340 $coord(id_1) = (20.88, 11.90)$ leaving object (X, Y) at T IFF walking(id₃) 340X is not inactive at T AND 340 $coord(id_3) = (24.78, 18.77)$ Y appears at T AND 380 walking(id₃) Y is inactive at T AND 380 $coord(id_3) = (27.88, 9.90)$ coord(X) equals coord(Y) at T 380 walking(id₂) 380 $coord(id_2) = (28.27, 9.66)$



53

Logic-based Event Recognition

- Formal semantics.
- Relational structure.
- Exploit background knowledge provided by domain experts.
- Reasoning about time and change using action formalisms.
- Event definitions may be learned from data.

The Event Calculus

Predicate	Description
HappensAt(e,t)	Event e occurs at time t
InitiatedAt(f,t)	Fluent f is initiated at time t
${\tt TerminatedAt}(f,t)$	Fluent f is terminated at time t
$\mathtt{HoldsAt}(f,t)$	Fluent f holds at time t

Axioms:

$$\begin{split} \operatorname{HoldsAt}(f,t+1) &\Leftarrow \\ \operatorname{InitiatedAt}(f,t) \\ \operatorname{HoldsAt}(f,t+1) &\Leftarrow \\ \operatorname{HoldsAt}(f,t) \land \\ \neg \operatorname{TerminatedAt}(f,t) \end{split}$$

 $\begin{array}{l} \neg \texttt{HoldsAt}(f,t{+}1) \Leftarrow \\ \texttt{TerminatedAt}(f,t) \\ \neg \texttt{HoldsAt}(f,t{+}1) \Leftarrow \\ \neg \texttt{HoldsAt}(f,t) \land \\ \neg \texttt{InitiatedAt}(f,t) \end{array}$



Learning Complex Event Definitions





ILP^1 and SRL^2 Systems

Input:

- Positive and negative examples.
- Background knowledge.
- Language bias.
- Output:
 - A logical theory that covers as many positive and as few negative examples as possible.
- Several methods exist for learning Event Calculus definitions:
 - ▶ XHAIL, ILED, OLED, OSL α , WOLED, ILASP.

Evangelos Michelioudakis

¹Inductive Logic Programming ²Statistical Relational Learning

Motivation: Semi-Supervised Learning of CE Definitions





Evangelos Michelioudakis

Overview of Semi-Supervised Learning



Scalable Semi-Supervised Structure Learning for Event Recognition

⁹/₅₃









Evangelos Michelioudakis





Evangelos Michelioudakis





Evangelos Michelioudakis

► Laplacian:

 $\mathbf{L} = \mathbf{D} - \mathbf{W} =$

0.5	0	0	0	- <mark>0</mark> .5	0	0	0	0]
0	0.95	0	0	0	- <mark>0</mark> .7	<mark>-0</mark> .25	0	0
0	0	0.9	0	0	0	0	0	<mark>-0</mark> .9
0	0	0	1.28	0	0	0	<mark>-0</mark> .71	<mark>-0</mark> .57
-0.5	0	0	0	1.4	-0.1	0	0	-0.8
0	- <mark>0</mark> .7	0	0	-0.1	0.8	0	0	0
0	<mark>-0</mark> .25	0	0	0	0	0.45	-0.67	-0.2
0	0	0	<mark>-0</mark> .71	0	0	-0.67	1.38	0
0	0	<mark>-0</mark> .9	<mark>-0</mark> .57	-0.8	0	-0.2	0	2.47

▶ The solution to the optimization is given by $\mathbf{y}_u = \mathbf{L}_{uu}^{-1} \mathbf{L}_{ul} \mathbf{y}_l$

Presentation Outline

1 Introduction

2 SPLICE: Semi-Supervised Learning for Complex Event Recognition

- 3 SPLICE⁺: Semi-Supervised Learning Combining Structure and Mass-based Predicate Similarity
- 4 Experimental Study
- 5 Conclusions & Future Work



Building on Graph-based SSL

Label Propagation:

- × Propositional approach that cannot learn logical theories.
- × Does not scale on large dataset.
- $\checkmark\,$ Its transductive nature enable it to be wrapped around ILP/SRL.

Goal:

- Infer the missing labels online (single-pass) using graph-based methods.
- Employ online ILP/SRL systems to learn the CE definitions.



Training Sequence

```
\begin{array}{l} \label{eq:micro-Batch $\mathcal{D}_t$} \\ \text{Happenskt}(walking(ID_1), 5) \\ \text{Happenskt}(walking(ID_2), 5) \\ \text{DrientationMove}(ID_1, ID_2, 5) \\ \text{Close}(ID_1, ID_2, 34, 5) \\ \text{Holdskt}(move(ID_1, ID_2), 5) \\ & \cdots \end{array}
```

```
\begin{array}{l} \mathtt{HappensAt}(\mathtt{exit}(\mathtt{ID}_1), 20) \\ \mathtt{HappensAt}(\mathtt{walking}(\mathtt{ID}_2), 20) \\ \neg \mathtt{OrientationMove}(\mathtt{ID}_1, \mathtt{ID}_2, 20) \\ \neg \mathtt{Close}(\mathtt{ID}_1, \mathtt{ID}_2, 34, 20) \\ \neg \mathtt{HoldsAt}(\mathtt{move}(\mathtt{ID}_1, \mathtt{ID}_2), 20) \end{array}
```

```
\label{eq:linear} \begin{split} & \texttt{HappensAt}(\texttt{walking}(\texttt{ID}_1),\texttt{50}) \\ & \texttt{HappensAt}(\texttt{walking}(\texttt{ID}_2),\texttt{50}) \\ & \texttt{OrientationMove}(\texttt{ID}_1,\texttt{ID}_2,\texttt{50}) \\ & \texttt{Close}(\texttt{ID}_1,\texttt{ID}_2,\texttt{34},\texttt{50}) \\ & \texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1,\texttt{ID}_2),\texttt{50}) \\ & \texttt{widdentationMove}(\texttt{ID}_1,\texttt{ID}_2),\texttt{50}) \end{split}
```

 $HoldsAt(move(ID_1, ID_2), 5)$

 $\neg \texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1,\texttt{ID}_2),\texttt{20})$

 $? HoldsAt(move(ID_1, ID_2), 50)$

PREDICATE SPECIFICATION: Happens&t(walking(person),time) Happens&t(exit(person),time) OrientationMove(person,person,time) Close(person,person),time) Holds&t(move(person,person),time)



Training Sequence

```
Micro-Batch D<sub>t</sub>
HappensAt(walking(ID<sub>1</sub>), 5)
HappensAt(walking(ID_2), 5)
OrientationMove(ID1, ID2, 5)
Close(ID1. ID2. 34.5)
HoldsAt(move(ID<sub>1</sub>, ID<sub>2</sub>), 5)
HappensAt(exit(ID1), 20)
HappensAt(walking(ID<sub>2</sub>), 20)
¬OrientationMove(ID1.ID2.20)
¬Close(ID1, ID2, 34, 20)
\neg HoldsAt(move(ID_1, ID_2), 20)
HappensAt(walking(ID_1), 50)
HappensAt(walking(ID_2), 50)
OrientationMove(ID1.ID2.50)
Close(ID1, ID2, 34, 50)
HoldsAt(move(ID, ID), 50)
```

 $HoldsAt(move(ID_1, ID_2), 5)$

 $\neg \texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1,\texttt{ID}_2),\texttt{20})$

 $? HoldsAt(move(ID_1, ID_2), 50)$

PREDICATE SPECIFICATION: HappensAt(walking(person), time) HappensAt(walking), time) OrientationMove(person, person, time) Close(person, person), time) HoldsAt(move(person, person), time)



Training Sequence

```
Micro-Batch D<sub>t</sub>
HappensAt(walking(ID1),5)
HappensAt(walking(ID_2), 5)
OrientationMove(ID1, ID2, 5)
Close(ID1, ID2, 34, 5)
HoldsAt(move(ID<sub>1</sub>, ID<sub>2</sub>), 5)
HappensAt(exit(ID1), 20)
HappensAt(walking(ID<sub>2</sub>), 20)
¬OrientationMove(ID1.ID2.20)
¬Close(ID1, ID2, 34, 20)
\neg HoldsAt(move(ID_1, ID_2), 20)
HappensAt(walking(ID_1), 50)
HappensAt(walking(ID_2), 50)
OrientationMove(ID1.ID2.50)
Close(ID1, ID2, 34, 50)
HoldsAt(move(ID, ID), 50)
```

 $HoldsAt(move(ID_1, ID_2), 5)$

 $\neg \texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1,\texttt{ID}_2),\texttt{20})$

 $? HoldsAt(move(ID_1, ID_2), 50)$

PREDICATE SPECIFICATION: HappensAt(walking(person), time) HappensAt(walking(person), time) OrientationMove(person, person, time) Close(person, person), time) HoldsAt(move(person, person), time)



Evangelos Michelioudakis

Training Sequence

```
Micro-Batch D<sub>t</sub>
HappensAt(walking(ID1),5)
HappensAt(walking(ID_2), 5)
OrientationMove(ID1, ID2, 5)
Close(ID1. ID2. 34.5)
HoldsAt(move(ID<sub>1</sub>, ID<sub>2</sub>), 5)
HappensAt(exit(ID1), 20)
HappensAt(walking(ID<sub>2</sub>), 20)
¬OrientationMove(ID1.ID2.20)
¬Close(ID1, ID2, 34, 20)
\neg HoldsAt(move(ID_1, ID_2), 20)
HappensAt(walking(ID_1), 50)
HappensAt(walking(ID_2), 50)
OrientationMove(ID1.ID2.50)
Close(ID1, ID2, 34, 50)
HoldsAt(move(ID, ID), 50)
```

 $\texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1\checkmark,\texttt{ID}_2),5)$

 $\neg \texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1\checkmark,\texttt{ID}_2),\texttt{20})$

?HoldsAt(move($ID_1 \checkmark$, ID_2), 50)

PREDICATE SPECIFICATION: HappensAt(walking(person), time) HappensAt(walking(person), time) OrientationMove(person, person, time) Close(person, person), time) HoldsAt(move(person, person), time)



Evangelos Michelioudakis

Training Sequence

```
Micro-Batch D<sub>t</sub>
HappensAt(walking(ID<sub>1</sub>), 5)
HappensAt(walking(ID_2), 5)
OrientationMove(ID1, ID2, 5)
Close(ID1, ID2, 34, 5)
HoldsAt(move(ID<sub>1</sub>, ID<sub>2</sub>), 5)
HappensAt(exit(ID1), 20)
HappensAt(walking(ID<sub>2</sub>), 20)
¬OrientationMove(ID1.ID2.20)
¬Close(ID1, ID2, 34, 20)
\neg HoldsAt(move(ID_1, ID_2), 20)
HappensAt(walking(ID_1), 50)
HappensAt(walking(ID_2), 50)
OrientationMove(ID1.ID2.50)
Close(ID1, ID2, 34, 50)
HoldsAt(move(ID, ID), 50)
```

HoldsAt(move(ID₁, ID₂), 5)

 $\neg \texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1,\texttt{ID}_2), \textcolor{red}{\texttt{20}})$

 $? HoldsAt(move(ID_1, ID_2), 50)$

PREDICATE SPECIFICATION: Happens&t(walking(person),time) Happens&t(exit(person),time) OrientationMove(person,person,time) Close(person,person),time) Holds&t(move(person,person),time)



Evangelos Michelioudakis

Training Sequence

```
Micro-Batch D<sub>t</sub>
HappensAt(walking(ID<sub>1</sub>), 5)
HappensAt(walking(ID_2), 5)
OrientationMove(ID1, ID2, 5)
Close(ID1. ID2. 34.5)
HoldsAt(move(ID<sub>1</sub>, ID<sub>2</sub>), 5)
HappensAt(exit(ID1), 20)
HappensAt(walking(ID<sub>2</sub>), 20)
¬OrientationMove(ID1.ID2.20)
¬Close(ID1, ID2, 34, 20)
\neg HoldsAt(move(ID_1, ID_2), 20)
HappensAt(walking(ID_1), 50)
HappensAt(walking(ID_2), 50)
OrientationMove(ID1.ID2.50)
Close(ID1, ID2, 34, 50)
HoldsAt(move(ID, ID), 50)
```

 $HoldsAt(move(ID_1, ID_2), 5\checkmark)$

 $HappensAt(walking(ID_1), 5)$

 $\neg \texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1,\texttt{ID}_2), \texttt{20} \times))$

 $? HoldsAt(move(ID_1, ID_2), 50 \times))$

PREDICATE SPECIFICATION: HappensAt(walking(person),time) HappensAt(wait(person),time) OrientationMove(person,person,time) Close(person,person,pixels,time) HoldsAt(move(person,person),time)

Training Sequence

```
Micro-Batch D<sub>t</sub>
HappensAt(walking(ID1),5)
HappensAt(walking(ID_2), 5)
OrientationMove(ID1, ID2, 5)
Close(ID1. ID2. 34.5)
HoldsAt(move(ID<sub>1</sub>, ID<sub>2</sub>), 5)
HappensAt(exit(ID1), 20)
HappensAt(walking(ID<sub>2</sub>), 20)
¬OrientationMove(ID1.ID2.20)
¬Close(ID1, ID2, 34, 20)
\neg HoldsAt(move(ID_1, ID_2), 20)
HappensAt(walking(ID_1), 50)
HappensAt(walking(ID_2), 50)
OrientationMove(ID1.ID2.50)
Close(ID1, ID2, 34, 50)
HoldsAt(move(ID, ID), 50)
```

 $HoldsAt(move(ID_1, ID_2 \checkmark), 5 \checkmark)$

HappensAt(walking(ID₁), 5) HappensAt(walking(ID₂), 5)

 $\neg \texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1, \texttt{ID}_2\checkmark), \texttt{20} \times))$

? HoldsAt(move($ID_1, ID_2\checkmark$), $50\times$))

PREDICATE SPECIFICATION: HappensAt(walking(person), time) HappensAt(walking(person), time) OrientationMove(person, person, time) Close(person, person, juels, time) HoldsAt(move(person, person), time)



Training Sequence



 $HoldsAt(move(ID_1, ID_2), 5)$

- ->

 $\begin{array}{l} \texttt{HappensAt}(\texttt{walking}(\texttt{ID}_1),\texttt{5})\\ \texttt{HappensAt}(\texttt{walking}(\texttt{ID}_2),\texttt{5})\\ \texttt{OrientationMove}(\texttt{ID}_1,\texttt{ID}_2,\texttt{5})\\ \texttt{Close}(\texttt{ID}_1,\texttt{ID}_2,\texttt{34},\texttt{5}) \end{array}$

 $\neg \texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1,\texttt{ID}_2),\texttt{20}))$

HappensAt(exit(ID₁), 20) HappensAt(walking(ID₂), 20)

 $? HoldsAt(move(ID_1, ID_2), 50))$

 $\begin{array}{l} \texttt{HappensAt}(\texttt{walking}(\texttt{ID}_1), \texttt{50}) \\ \texttt{HappensAt}(\texttt{walking}(\texttt{ID}_2), \texttt{50}) \\ \texttt{OrientationMove}(\texttt{ID}_1, \texttt{ID}_2, \texttt{50}) \\ \texttt{Close}(\texttt{ID}_1, \texttt{ID}_2, \texttt{34}, \texttt{50}) \end{array}$

PREDICATE SPECIFICATION: HappensAt(walking(person), time) HappensAt(walking(person), time) OrientationMove(person, person, time) Close(person, person, time) HoldsAt(move(person, person), time)



Label Caching

Store labelled examples for future usage:







Evangelos Michelioudakis

Label Caching



Evangelos Michelioudakis

Label Caching





The Hoeffding Bound

- X is a random variable.
- X_1, \ldots, X_N are N independent observation of X values.
- Let \overline{X} be the known, observed mean of X.
- Let \hat{X} be the unknown, true mean of X.
- ► Then, it holds that:

$$ar{X} - \epsilon \leq \hat{X} \leq ar{X} + \epsilon$$
, with probability, $1 - \delta$, where $\epsilon = \sqrt{rac{\ln(1/\delta)}{2N}}$



Monitor Contradicting Examples

- Let p_c be the observed probability of observing labelled example c.
- Let $p_{\neg c}$ be the observed probability of observing the opposite example.
- As examples stream-in, monitor the quantity $\bar{X} = p_c p_{\neg c}$.
- Continue until N examples makes $\bar{X} > \epsilon$.



Graph Construction

 $\begin{array}{l} \texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1,\texttt{ID}_2),\texttt{5})\\ \texttt{HappensAt}(\texttt{walking}(\texttt{ID}_1),\texttt{5})\\ \texttt{HappensAt}(\texttt{walking}(\texttt{ID}_2),\texttt{5})\\ \texttt{OrientationMove}(\texttt{ID}_1,\texttt{ID}_2,\texttt{5})\\ \texttt{Close}(\texttt{ID}_1,\texttt{ID}_2,\texttt{34},\texttt{5}) \end{array}$

$$\label{eq:holdsAt} \begin{split} &\neg \texttt{HoldsAt}(\texttt{move}(\texttt{ID}_1,\texttt{ID}_2),\texttt{20}) \\ &\texttt{HappensAt}(\texttt{exit}(\texttt{ID}_1),\texttt{20}) \\ &\texttt{HappensAt}(\texttt{walking}(\texttt{ID}_2),\texttt{20}) \end{split}$$

 $\label{eq:holdsAt} \begin{array}{l} ? \mbox{HoldsAt}(\mbox{move}(ID_1, ID_2), 50) \\ \mbox{HappensAt}(\mbox{walking}(ID_1), 50) \\ \mbox{HappensAt}(\mbox{walking}(ID_2), 50) \\ \mbox{OrientationMove}(ID_1, ID_2, 50) \\ \mbox{Close}(ID_1, ID_2, 4, 50) \end{array}$



Evangelos Michelioudakis

Nienhuys-Cheng (NC) Distance

- Distance for Herbrand interpretations.
- Let \mathcal{E} be the set of all expressions in a first-order language.
- The distance $d : \mathcal{E} \times \mathcal{E} \mapsto \mathbb{R}$ is defined as follows:

$$d(e, e) = 0, \forall e \in \mathcal{E}$$

$$d(p(s_1, \dots, s_k), q(t_1, \dots, t_r)) = 1, p \neq q \lor k \neq r$$

$$d(p(s_1, \dots, s_k), q(t_1, \dots, t_k)) = \frac{1}{2k} \sum_{i=1}^k d(s_i, t_i), p = q$$



Nienhuys-Cheng (NC) Distance

The distance between evidence atoms can be computed as follows:

- $e_1 = \texttt{HappensAt}(\texttt{walking}(\texttt{ID}_1), 5)$
- $\blacktriangleright \ e_2 = \texttt{HappensAt}(\texttt{walking}(\texttt{ID}_2), \texttt{50})$

$$d(e_1, e_2) = \frac{1}{2p} \sum_{i=1}^p d(e_{1,p}, e_{2,p}) = \frac{1}{2 \cdot 2} \Big(d(\operatorname{walking}(\operatorname{ID}_1), \operatorname{walking}(\operatorname{ID}_2)) + d(5, 50) \Big) = \frac{1}{4} \Big(\frac{1}{2 \cdot 1} d(\operatorname{ID}_1, \operatorname{ID}_2) + 1 \Big) = \frac{1}{4} \Big(\frac{1}{2} \cdot 1 + 1 \big) = 0.375$$



Kuhn-Munkres Algorithm

- Given a M × M non-negative cost matrix C, find a one-to-one mapping from rows to columns, such that the total cost is minimum.
- Assuming a permutation matrix P, the optimization problem may be formulated as follows:

 $\min_{\mathbf{P}} \operatorname{Tr}(\mathbf{PC})$

In the general case, of a M × K matrix, where M > K, C is padded to complete the smaller dimension.


Kuhn-Munkres distance for Herbrand Interpretations

Given a pair evidence atom sets \mathcal{E}_i and \mathcal{E}_j , the cost matrix is computed using the Nienhuys-Cheng distance:

$$\mathbf{C} = \begin{pmatrix} d(e_{i,1}, e_{j,1}) & d(e_{i,1}, e_{j,2}) & \cdots & d(e_{i,1}, e_{j,M}) \\ d(e_{i,2}, e_{j,1}) & d(e_{i,2}, e_{j,2}) & \cdots & d(e_{i,2}, e_{j,M}) \\ \vdots & \vdots & \ddots & \vdots \\ d(e_{i,M}, e_{j,1}) & d(e_{i,M}, e_{j,2}) & \cdots & d(e_{i,M}, e_{j,M}) \end{pmatrix}$$

The proposed distance is defined as follows:

$$d_s(\mathcal{E}_i, \mathcal{E}_j) = \frac{1}{M} \operatorname{Tr}(\mathbf{P}^* \mathbf{C})$$



Kuhn-Munkres distance for Herbrand Interpretations



Then, the following permuation matrix gives the minimum total cost:

$$\mathbf{P}^{\star} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$



Kuhn-Munkres distance for Herbrand Interpretations



$$\operatorname{Tr}(\mathbf{P}^{\star}\mathbf{C}) = \frac{1}{4}(0.25 + 1 + 0.25 + 1) = 0.75$$

• The runtime complexity for finding \mathbf{P}^{\star} is $\mathcal{O}(\max(|\mathcal{E}_i|, |\mathcal{E}_j|)^3)$



Graph Construction



$$\mathbf{W} = \begin{bmatrix} 0 & 0.81 & 0.81 \\ 0.81 & 0 & 0.19 \\ 0.81 & 0.19 & 0 \end{bmatrix}$$



Supervision Completion through Label Propagation

Given the adjacency matrix $\mathbf W$, we compute the degree matrix,

$$\mathbf{W} = \begin{bmatrix} 0 & 0.81 & 0.81 \\ 0.81 & 0 & 0.19 \\ 0.81 & 0.19 & 0 \end{bmatrix} \longrightarrow \mathbf{D} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1.62 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

and the Laplacian matrix:

$$\mathbf{L} = \mathbf{D} - \mathbf{W} = \begin{bmatrix} 1 & -0.81 & -0.91 \\ -0.81 & 1.62 & -0.81 \\ -0.19 & -0.81 & 1 \end{bmatrix}$$

The solution for the unlabelled examples is obtained as follows:

$$\mathbf{f}_u = \mathbf{L}_{uu}^{-1} \mathbf{L}_{ul} \mathbf{f}_l = \begin{bmatrix} 0.62 \end{bmatrix}$$

${\rm Splice} \ Overview$





- $\checkmark\,$ Facilitates learning CE rules in the presence of incomplete supervision.
- ✓ Handles contradicting annotations.
- $\checkmark\,$ Scales to large datasets.
- imes The distance measure is deluded in the presence of irrelevant features.
- \times The online labelling is suboptimal.

Presentation Outline

1 Introduction

- 2 SPLICE: Semi-Supervised Learning for Complex Event Recognition
- 3 SPLICE⁺: Semi-Supervised Learning Combining Structure and Mass-based Predicate Similarity
- 4 Experimental Study
- 5 Conclusions & Future Work



Large Margin Nearest Neighbour (LMNN) Metric Learning





Large-Margin Feature Selection for Logical Predicates

- Adapt LMNN metric learning for feature selection.
- Let A_l be a totally ordered set of atoms constructed by:
 - Hebrand base \mathcal{B}_l defined by the labelled instances.
 - Language bias.

• Each labelled instance is represented as a binary vector $\mathbf{x} = [x_1, \dots, x_{|\mathcal{A}_l|}]^{\mathsf{T}}$.





Large-Margin Feature Selection for Logical Predicates

- ► Distance between vectors x can be defined as a Hamming distance.
- Reformulate the cost function of LMNN:

$$\varepsilon(\mathbf{b}) = \sum_{i,j\in\mathcal{N}_i^k} \sum_{l} (1-y_{il}) \left[1 + \mathbf{b} |\mathbf{x}_i - \mathbf{x}_j| - \mathbf{b} |\mathbf{x}_i - \mathbf{x}_l| \right]_+$$

- The binary vector **b** denotes which features in A_l are selected.
- Since b is a vector, instead of a matrix, the problem can be solved using integer programming.



Mass Distance for Logical Predicates

LMFS require labelled data for performing feature selection.

- The available labels may not be sufficient to find all irrelevant features.
- Employ a data-driven metric based on mass estimation.
 - Examples are considered more similar if they coexist in a sparse space.





Mass Distance for Logical Predicates

• Create hierarchical partitionings of the Hebrand base \mathcal{B}_{l+u} .

- Half-Space trees (represent part of the subsumption lattice).
- Trees can be constructed beforehand.
- Update the mass of each subspace from incoming labelled and unlabelled examples.
- ▶ Estimate the distance (mass) from a set of *T* Half-Space trees:

$$\tilde{m}(\mathbf{x}, \mathbf{y}) = \frac{1}{T} \sum_{p=1}^{T} \frac{|R(\mathbf{x}, \mathbf{y}|H_p)|}{|D|}$$



Mass Distance for Logical Predicates



Scalable Semi-Supervised Structure Learning for Event Recognition

53

Edge similarities are measured using the combined measure:

$$d_h^{\mathbf{b}}(v_i, v_j) = \alpha \ d_s^{\mathbf{b}}(v_i, v_j) + (1 - \alpha) \ \tilde{m}(v_i, v_j)$$

▶ Vertices are connected using a temporal variant of *k*NN.

- Connects unlabelled examples to their k-nearest labelled neighbours.
- Connects unlabelled examples to their temporally adjacent ones.
- Maintains a synopsis graph S that contains:
 - The τ most recent unlabelled examples.
 - All cached labelled examples.
- Removes older than τ unlabelled examples using star-mesh transforms.
 - The labeling solution on \mathcal{G} is preserved on \mathcal{S} .

Training sequence:





Evangelos Michelioudakis

Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 0$):





Evangelos Michelioudakis

Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 0$):







Evangelos Michelioudakis

Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 1$):





Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 1$):





Evangelos Michelioudakis

Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 1$):





Evangelos Michelioudakis

Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 2$):





Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 2$):





Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 2$):





Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 3$):





Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 3$):





Training sequence:

Graph synopsis \mathcal{S} (memory $\tau = 2$):





Evangelos Michelioudakis

Training sequence:

Graph synopsis \mathcal{S} (memory $\tau = 2$):





Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 3$):





Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 3$):





Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 2$):





Training sequence:



Graph synopsis \mathcal{S} (memory $\tau = 2$):





Training sequence:







Training sequence:





Training sequence:







Training sequence:



Graph synopsis S (memory $\tau = 2$): $w_{v,v'} = w_{v,v'} + \frac{w_{v,v_0}w_{v_0,v'}}{2.45}$



Evangelos Michelioudakis
Presentation Outline

1 Introduction

- 2 SPLICE: Semi-Supervised Learning for Complex Event Recognition
- 3 SPLICE⁺: Semi-Supervised Learning Combining Structure and Mass-based Predicate Similarity
- 4 Experimental Study
- 5 Conclusions & Future Work



Datasets

Task	CE	Timestamps	Occurrences
Activity	meet	12,869	6,272
Recognition	move	$12,\!869$	3,722
Maritime	rendezVous	11,930	1,425
Monitoring	pilotOps	$6,\!678$	769
Fleet	nonEconomicDriving	$13,\!255$	1,589
Management	dangerousDriving	$13,\!387$	639



Experimental Setup

Scenario I: Labelled micro-batches were randomly selected.

- \blacktriangleright Retaining 5%, 10%, 20%, 40% and 80% of the micro-batches labelled.
- ▶ The random selection was repeated 20 times for each percentage.
- Scenario II: Labelled micro-batches appear at the beginning of training.
 - Each dataset comprise a set of independant sequences.
 - Each sequence was selected to appear as fully labelled, while others followed completely unlabelled.



Activity Recognition: Supervision Completion



Scalable Semi-Supervised Structure Learning for Event Recognition

/53

Activity Recognition: Runtime Performance



Scalable Semi-Supervised Structure Learning for Event Recognition

⁴²/₅₃

Activity Recognition: Structure Learning



Scalable Semi-Supervised Structure Learning for Event Recognition

 $^{43}/_{53}$

Maritime Monitoring: Supervision Completion



Scalable Semi-Supervised Structure Learning for Event Recognition

/53

Maritime Monitoring: Structure Learning



Scalable Semi-Supervised Structure Learning for Event Recognition

⁴⁵/₅₃

Fleet Management: Supervision Completion



Scalable Semi-Supervised Structure Learning for Event Recognition

53

Fleet Management: Structure Learning



Scalable Semi-Supervised Structure Learning for Event Recognition

/53

Batch Size Impact (SPLICE vs SPLICE⁺)

CE	Batch size	Splice/Splice ⁺		
		Number of supervised sequences		
		1	2	4
pilotOps	10	0.63/0.96	0.88/0.97	0.92/0.97
	25	0.69/0.96	0.85/0.96	0.91/0.97
	50	0.71/0.96	0.88/0.96	0.91/0.97
	100	0.61/0.95	0.70/0.96	0.75/0.97
rendezVous	10	0.63/0.74	0.77/0.86	0.87/0.93
	25	0.58/0.74	0.72/0.86	0.84/0.90
	50	0.56/0.74	0.75/0.86	0.83/0.90
	100	0.48/0.75	0.61/0.81	0.83/0.92



Evangelos Michelioudakis

Presentation Outline

1 Introduction

- 2 SPLICE: Semi-Supervised Learning for Complex Event Recognition
- 3 SPLICE⁺: Semi-Supervised Learning Combining Structure and Mass-based Predicate Similarity
- 4 Experimental Study
- 5 Conclusions & Future Work



Conclusions

We presented two scalable algorithms for semi-supervised learning of complex event definitions in the form of Event Calculus theories 3 .

► Splice

- Online learner.
- Infers missing labels using label propagation and a distance measure for Herbrand Interpretations.
- Caches previously seen labelled examples and detects contradictions.
- ► Splice⁺
 - Accounts for irrelevant or noisy features in distance measurements.
 - Structural distance optimised through feature selection.
 - Data-driven distance based on mass estimation.
 - Maintains a synopsis of the input examples to achieve optimal labelling.

³https://github.com/anskarl/LoMRF

Publications

Primary

- Michelioudakis, E., Artikis, A. and Paliouras, G. (2023), "Online Semi-Supervised Learning of Event Rules Combining Structure and Mass-Based Predicate Similarity", *Machine Learning (accepted)*
- Michelioudakis, E., Artikis, A. and Paliouras, G. (2019), "Semi-Supervised Online Structure Learning for Composite Event Recognition", *Machine Learning*, 108(7), pp. 1085–1110

Secondary

- Akasiadis, C., Ponce-de-Leon, M., Montagud, A., Michelioudakis, E., Atsidakou, A., Alevizos, E., Artikis, A., Valencia, A., Paliouras G. (2022), "Parallel Model Exploration for Tumor Treatment Simulations", *Computational Intelligence*, 38(4), pp. 1379–1401
- Stavropoulos, V., <u>Michelioudakis, E.</u>, Akasiadis, C., Artikis, A. (2022), "Resource-Effective Exploration of Tumor Treatments with Multi-scale Simulations", *Hellenic Conference on Artificial Intelligence*, pp. 1–10

Katzouris, N., <u>Michelioudakis, E.</u>, Artikis, A., Paliouras, G. (2018), "Online Learning of Weighted Relational Rules for Complex Event Recognition", In Proceedings of the ECML-PKDD, pp. 396–413

Future Work

- Active learning.
 - Select unlabelled instances and request feedback.
 - Detect regions of the input stream where SDEs change significantly.
 - Probabilistic inference and uncertainty sampling.
- Distributed learning.
 - Inferring the missing labels for different CEs is trivially parallel.
 - Partitioning the data based on the constants appearing in the CEs.



Thank you!



Evangelos Michelioudakis

Appendix



Evangelos Michelioudakis

Online Learning of Event Definitions (OLED)



Scalable Semi-Supervised Structure Learning for Event Recognition

53

Shur's Complement

- Let $V = V_a \cup V_b$ be a partition of the vertices in a graph \mathcal{G} .
- The block partition of the Laplacian is as follows:

$$\mathbf{L} = egin{bmatrix} \mathbf{L}_{aa} & \mathbf{L}_{ab} \ \mathbf{L}_{ba} & \mathbf{L}_{bb} \end{bmatrix}$$

The Shur's Complement states that for the short-circuit graph defined on the vertices V_b the Laplacian is as follows:

$$\mathbf{L}/\mathbf{L}_{aa} = \mathbf{L}_{bb} - \mathbf{L}_{ba}\mathbf{L}_{aa}^{-1}\mathbf{L}_{ab}$$

The reduced graph can be computed by sequential star-mesh transforms on the V_a vertices.

A star-mesh transform on a vertex v_o of a given graph $\mathcal{G}=(V, E, \mathbf{W})$ is defined as follows:

- 1. Remove v_o from \mathcal{G} together with its set E_o of incident edges $(v_o, v) \in E_o$.
- 2. For every pair of vertices $v, v' \in V$ such that $(v_o, v) \in E_o$ and $(v_o, v') \in E_o$, add the edge (v, v') to E with weight $w_{v,v'} = w_{v,v_o} w_{v_o,v'} / degree(v_o)$. If (v, v') is already in E, then add the new weight $w_{v,v'}$ to its current weight.



- Nienhuys-Cheng distance employs the Hausdorff metric to compute the distance between sets of atoms.
- Given the sets \mathcal{E}_1 and \mathcal{E}_2 , their distance is computed as follows:

 $\max\{\sup_{x\in\mathcal{E}_1}\inf_{y\in\mathcal{E}_2}d(x,y),\sup_{y\in\mathcal{E}_2}\inf_{x\in\mathcal{E}_1}d(x,y)\}$

Determined by the most distant elements of the sets, chosen among the nearest neighbors on the other set.





$$D = \begin{array}{ccc} e_{21} & e_{22} & e_{23} \\ e_{11} & 0.25 & 1 & 1 \\ 0.375 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 0.25 \end{array}$$







Evangelos Michelioudakis









Evangelos Michelioudakis



$$D = \begin{array}{ccc} e_{21} & e_{22} & e_{23} \\ e_{11} & 0.25 & 1 & 1 \\ 0.375 & 1 & 1 \\ e_{13} & e_{14} \end{array} \begin{pmatrix} 0.25 & 1 & 1 \\ 0.375 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 0.25 \end{pmatrix}$$



Evangelos Michelioudakis



$$D = \begin{array}{ccc} e_{21} & e_{22} & e_{23} \\ e_{11} & \begin{pmatrix} 0.25 & 1 & 1 \\ 0.375 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 0.25 \end{pmatrix}$$





$$D = \begin{array}{ccc} e_{21} & e_{22} & e_{23} \\ e_{11} & \begin{pmatrix} 0.25 & 1 & 1 \\ 0.375 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 0.25 \end{pmatrix}$$



Evangelos Michelioudakis



$$D = \begin{array}{ccc} e_{21} & e_{22} & e_{23} \\ e_{11} & 0.25 & 1 & 1 \\ 0.375 & 1 & 1 \\ e_{13} & e_{14} & 1 & 1 \\ 1 & 1 & 1 & 0.25 \end{array}$$





$$D = \begin{array}{ccc} e_{21} & e_{22} & e_{23} \\ e_{11} & 0.25 & 1 & 1 \\ 0.375 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 0.25 \end{array}$$

Drawbacks:

- Does not capture much information about the sets being compared.
- Not representative of the overall dissimilarity between clauses.

Kuhn-Munkres Algorithm

- 1. Subtract the smallest entry in each row from the row.
- 2. Subtract the smallest entry in each column from the column.
- 3. Draw the fewest row/column lines through the 0 entries.
- 4. If N lines drawn, an optimal assignment of zeros is possible, otherwise,
- 5. Find the smallest entry not covered by any line. Subtract this entry from each row that isn't crossed out, and then add it to each column that is crossed out. Go back to step 3.



Large Margin Nearest Neighbour (LMNN) Metric Learning

- ► Learns a Mahalannobis distance optimized for kNN classification. $d_{\mathbf{M}}(\mathbf{x}_i, \mathbf{x}_i) = (\mathbf{x}_i - \mathbf{x}_i)^{\mathsf{T}} \mathbf{M}(\mathbf{x}_i - \mathbf{x}_i)$
- Minimises a loss function consisting of two terms:
 - One that penalises large distances between instances having same labels.
 - One that penalises small distances between differently labelled examples.

$$\epsilon(\mathbf{M}) = (1-\mu) \sum_{i,j \in \mathcal{N}_i^k} d_{\mathbf{M}}(\mathbf{x}_i, \mathbf{x}_j) + \mu \sum_{i,j \in \mathcal{N}_i^k} \sum_l (1-y_{il}) [1 + d_{\mathbf{M}}(\mathbf{x}_i, \mathbf{x}_j) - d_{\mathbf{M}}(\mathbf{x}_i, \mathbf{x}_l)]_+$$



Large-Margin Feature Selection for Logical Predicates

Solve the following binary integer programming problem:

$$\begin{array}{ll} \text{minimise} & \sum_{i,j\in\mathcal{N}_i^k}\sum_l (1-y_{il})\xi_{ijl}\\ \text{subject to} & \textbf{(1) } \mathbf{b}|\mathbf{x}_i - \mathbf{x}_l| - \mathbf{b}|\mathbf{x}_i - \mathbf{x}_j| \geq 1 - \xi_{ijl}\\ & \textbf{(2) } \mathbf{b}\mathbf{x}_i \geq 1\\ & \textbf{(3) } \xi_{ijl} \in \mathbb{N}^{\geq}\\ & \textbf{(4) } \mathbf{b} \in \{0,1\}^{|\mathcal{A}|} \end{array}$$

Given the feature subset $\mathbf{b},$ the structural distance is computed as follows:

$$d_s^{\mathbf{b}}(\mathcal{E}_i, \mathcal{E}_j) = d_s(\mathcal{E}_i^{\mathbf{b}}, \mathcal{E}_j^{\mathbf{b}})$$



Evangelos Michelioudakis

Mass-based Distance

- The distance depends on the amount of probability mass in the region of space covering the two points.
- Let H denote a hierarchical partitioning of a space \mathbb{R}^q .
 - A set of non-overlapping regions that collectively span \mathbb{R}^q .
 - Each region in the *H* corresponds to the union of its child regions.
- The smallest region covering a pair \mathbf{x}, \mathbf{y} in H is defined as:

$$R(\mathbf{x}, \mathbf{y}|H) = \operatorname*{argmax}_{r \in H \ s.t. \{\mathbf{x}, \mathbf{y}\} \in r} \operatorname{depth}(r; H)$$

Assuming an unknown probability distribution F, the distance of x, y is defined as the expectation of a randomly chosen point to lie in R:

$$m(\mathbf{x}, \mathbf{y}|D) = E_{\mathcal{H}(D)} [P_F(R(\mathbf{x}, \mathbf{y}|H; D))]$$

• Estimates the mass from a set of Isolation Forest partitions over \mathbb{R}^q :

$$\tilde{m}(\mathbf{x}, \mathbf{y}) = \frac{1}{T} \sum_{p=1}^{T} \frac{|R(\mathbf{x}, \mathbf{y}|H_p)|}{|D|}$$

Experimental Setup





Evangelos Michelioudakis

Maritime Monitoring: Runtime Performance



Scalable Semi-Supervised Structure Learning for Event Recognition

⁶⁵/₅₃
Fleet Management: Runtime Performance



Scalable Semi-Supervised Structure Learning for Event Recognition

⁶⁶/₅₃

Batch Size Impact (SPLICE vs SPLICE⁺)

CE	Batch size	Number of supervised sequences				
		1	2	4	8	
meet	10	0.44/0.69	0.59/0.78	0.73/0.78	0.78/0.93	
	25	0.43/0.69	0.57/0.74	0.72/0.78	0.78/0.93	
	50	0.42/0.69	0.51/0.77	0.67/0.77	0.77/0.93	
	100	0.42/0.69	0.56/0.76	0.75/0.80	0.77/0.93	
move	10	0.66/0.73	0.73/0.75	0.71/0.79	0.84/0.94	
	25	0.66/0.73	0.74/0.74	0.72/0.79	0.84/0.94	
	50	0.66/0.73	0.74/0.78	0.74/0.81	0.84/0.94	
	100	0.66/0.73	0.73/0.75	0.73/0.80	0.84/0.94	



Ablation Study: Activity recognition

CE	Distance	Random Supervision		Early Supervision	
		5%	10%	1	2
meet	d_s	0.62	0.70	0.56	0.69
	$d_s^{\mathbf{b}}$	0.59	0.70	0.64	0.71
	m	0.65	0.75	0.64	0.70
	d_h	0.63	0.76	0.65	0.73
	$d_h^{\mathbf{b}}$	0.67	0.77	0.70	0.76
move	d_s	0.57	0.64	0.67	0.71
	$d_s^{\mathbf{b}}$	0.58	0.67	0.69	0.71
	m	0.56	0.68	0.60	0.54
	d_h	0.57	0.69	0.70	0.73
	$d_h^{\mathbf{b}}$	0.58	0.69	0.73	0.75



Evangelos Michelioudakis

Scalable Semi-Supervised Structure Learning for Event Recognition

Ablation Study: Maritime Monitoring

CE	Distance	Random Supervision		Early Supervision	
CE		5%	10%	1	2
	d_s	0.59	0.70	0.69	0.79
	$d_s^{\mathbf{b}}$	0.59	0.70	0.69	0.79
rendezVous	m	0.53	0.65	0.53	0.53
	d_h	0.62	0.75	0.74	0.81
	$d_h^{\mathbf{b}}$	0.62	0.75	0.74	0.81
	d_s	0.47	0.63	0.78	0.90
	$d_s^{\mathbf{b}}$	0.47	0.63	0.78	0.90
pilotOps	m	0.56	0.69	0.94	0.94
	d_h	0.56	0.69	0.95	0.96
	$d_h^{\mathbf{b}}$	0.56	0.69	0.95	0.96

