Online Learning of Weighted Relational Rules for Complex Event Recognition

Nikos Katzouris[1], Alexander Artikis[2,1] and Georios Paliouras[1]

http://cer.iit.demokritos.gr

[1]National Center for Scientific Research Demokritos, Athens, Greece
[2]University of piraeus, piraeus, Greece

ECML-PKDD 2018

The problem Setting

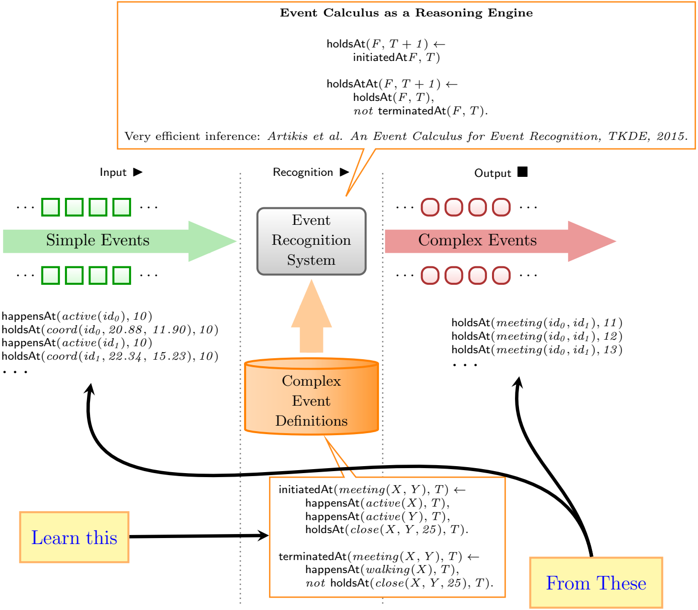

Learning Requirements

- Event recognition applications deal with noisy data streams.
  - Resilience to noise $\rightarrow$ Statistical Relational Learning.
  - Learning should be online.
    - Single-pass.
    - Learn from past mistakes.

Contribution of this Work

Two online learners from previous work:

- OLED
  - Katzouris N. et al. Online Learning of Event Definitions, *TPLP*, 2016.
  - ✓ Efficient structure learning using Hoeffding bounds.
  - ✗ Crisp learner.
- OSLα
  - Micheloudakis V., et al. OSLa: Online Structure Learning using Background Knowledge Axiomatization, *ECML*, 2016.
  - MLN learner.
  - ✓ Efficient weight learning.
  - ✗ Inefficient structure learning.
    - Blindly generates too many rules.

Current work:

- WoLED (OLED + weight learning)
  - MLN learner
  - ✓ Efficient structure learning.
  - ✓ Efficient weight learning.

OLED

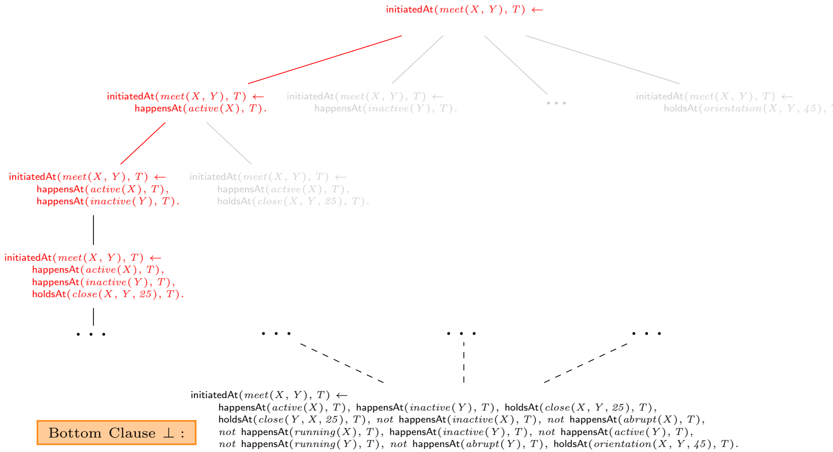

- Learns a rule with online hill-climbing.

OLED

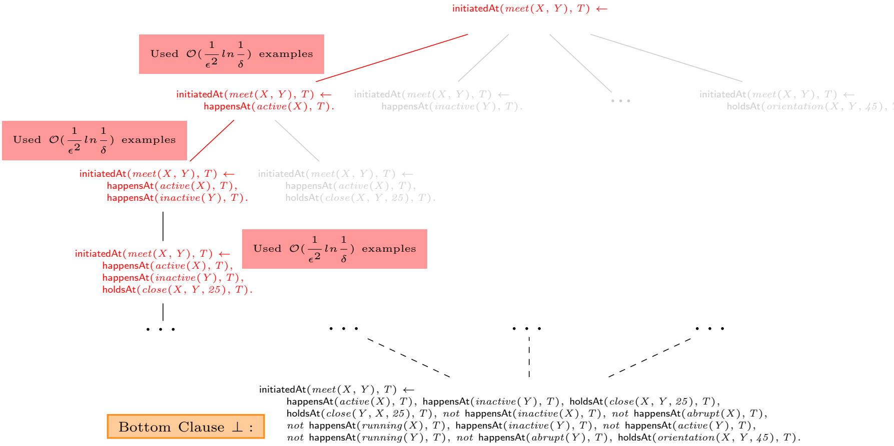

- Uses Hoeffding tests to make (ϵ, δ)-optimal decisions.

WoLED

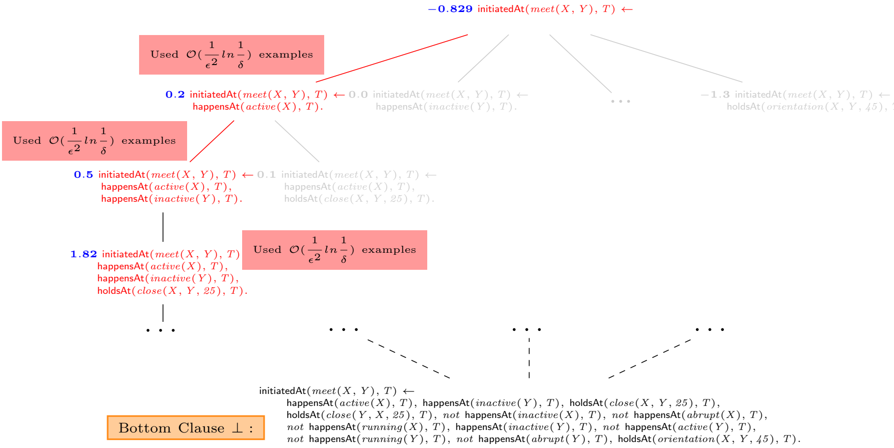

- The WoLED algorithm:
  - Simultaneous structure & weight learning.
  - Weight learning with AdaGrad.

The AdaGrad Weight Update Rule

Previous weight of the i-th rule

Learning rate

Rule's current mistakes

Regularization rate

$$w_i^{t+1} = sign(w_i^t - \frac{\eta}{C_i^t}\Delta g_i^t)\ max\{0, |w_i^t - \frac{\eta}{C_i^t}\Delta g_i^t| - \lambda\frac{\eta}{C_i^t}\}$$

Current weight of the i-th rule

Term proportional to the rule's accumulated past mistakes

- Δg_i^t (i-th rule's mistakes at time t): difference in rule's true groundings in the true state and the MAP-inferred state.

WoLED Overview

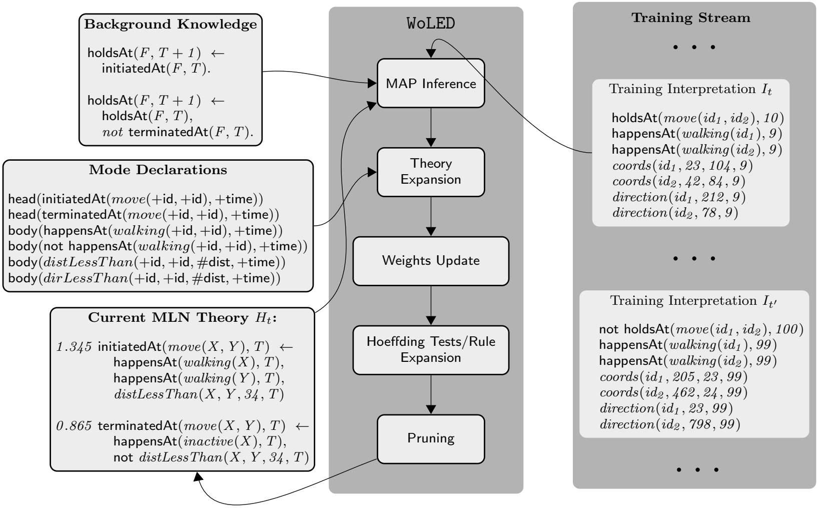

WoLED Evaluation on the CAVIAR Dataset

		Method	Precision	Recall	F_1-score	Theory size	Time (sec)
(a)	*Moving*	EC_{crisp}	**0.909**	0.634	0.751	28	–
		OLED	0.867	0.724	0.789	34	**28**
		WoLED	0.882	0.835	0.857	30	59
		OSLα	0.837	0.590	0.692	3316	1300
		OSL	-	-	-	-	$>$ 25 hrs
		MaxMargin	0.844	**0.941**	**0.890**	28	1692
		XHAIL	0.779	0.914	0.841	**14**	7836
	Meeting	EC_{crisp}	0.687	0.855	0.762	23	–
		OLED	**0.947**	0.760	0.843	31	**22**
		WoLED	0.892	0.888	**0.889**	29	52
		OSLα	0.902	0.863	0.882	1231	180
		OSL	-	-	-	-	$>$ 25 hrs
		MaxMargin	0.919	0.813	0.863	23	1133
		XHAIL	0.804	**0.927**	0.861	**15**	7248
(b)	*Moving*	OLED	0.682	0.787	0.730	**38**	**63**
		WoLED	**0.783**	**0.821**	**0.801**	51	108
		EC_{crisp}	0.721	0.639	0.677	28	–
	Meeting	OLED	0.701	0.886	0.782	**41**	**43**
		WoLED	**0.808**	**0.877**	**0.841**	56	98
		EC_{crisp}	0.644	0.855	0.735	23	–

Table 1: Experimental results on (a) a fragment of CAVIAR (top) and (b) the complete CAVIAR dataset (bottom).

WoLED Holdout Evaluation

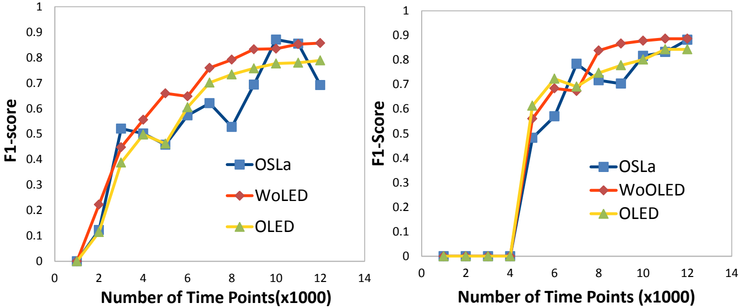

Summary

- An efficient, online MLN learner (structure+weights).
- Built on top of the LoMRF[1] platform.
- https://github.com/nkatzz/OLED

Future work:

- Further evaluation.
- Different Weight learning schemes.
- Distributed learning.

[1] https://github.com/anskarl/LoMRF