Events, Relationships, and Script Learning

Trigger Identification:

Bi-directional LSTM + CRF

- Concatenate character embeddings with pre-trained word embeddings to get vectors representing each word
- Run a bi-LSTM over the sequence of word vectors to obtain the two hidden states
- Use a CRF to find the sequence with the highest probability

Results	P	R	F1
JointEventEntity	61.4	71.0	65.9
Bi-LSTM + CRF	66.7	66.0	**66.4**

- Train SVMs to label triggers as belonging to one of 33 subtypes (e.g., attack, sentence, convict, etc.)
- However, with perfect trigger identification, F1 for trigger classification is 80.0.

Results	P	R	F1
JointEventEntity (reported)	75.1	63.3	68.7
JointEventEntity (reported)	61.5	71.2	66.0
SVM	81.6	54.5	65.3

Bottom line: Neither technique should be trusted to perform accurate trigger identification

Macro-Event Extraction

Document-level analysis – extract only the most important events in a text

The Shield Actor **Michael Jace** Found Guilty of Wife's Murder, Sentenced to 40 Years in Prison

ATTACK Macro-Event	
Perpetrator	Michael Jace
Victim – Dead	April Jace
Victim – Injured	None
Time	May 19, 2014
Location	Los Angeles

ARREST Macro-Event	
Arrestee	Michael Jace
Time	2014
Location	Los Angeles

TRIAL Macro-Event	
Defendant	Michael Jace
Crime	Murder
Verdict	Guilty
Sentence	Prison
Time	Tuesday
Location	Los Angeles

Developing two novel ML-based algorithms for solving this problem

- A structured prediction model based on Learning to Search
- A deep neural network based on machine comprehension with no reliance on target domain training data

Preliminary results on the attack and elections domains show significantly improved performance against baseline methods

Currently gathering annotated data via Mechanical Turk

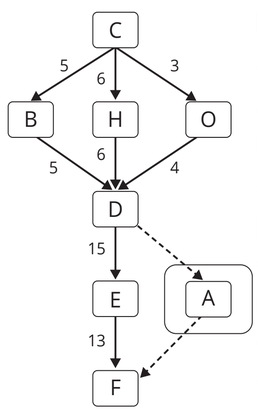

Extensions to FY16 Script Learning

- Better representation of the world
- Finer-grain event representation (actor, location, etc.)
- Probability distributions over the possible arguments
- A refined similarity metric that can reflect the new representation
- More robust script manipulation
- Better script addition (avoiding adding rare instances, etc.)
- Splitting / pruning existing scripts

Use of macro-event knowledge for inferring constraints

Learning to Extract Events Minimal Supervision

Gun Violence Archive

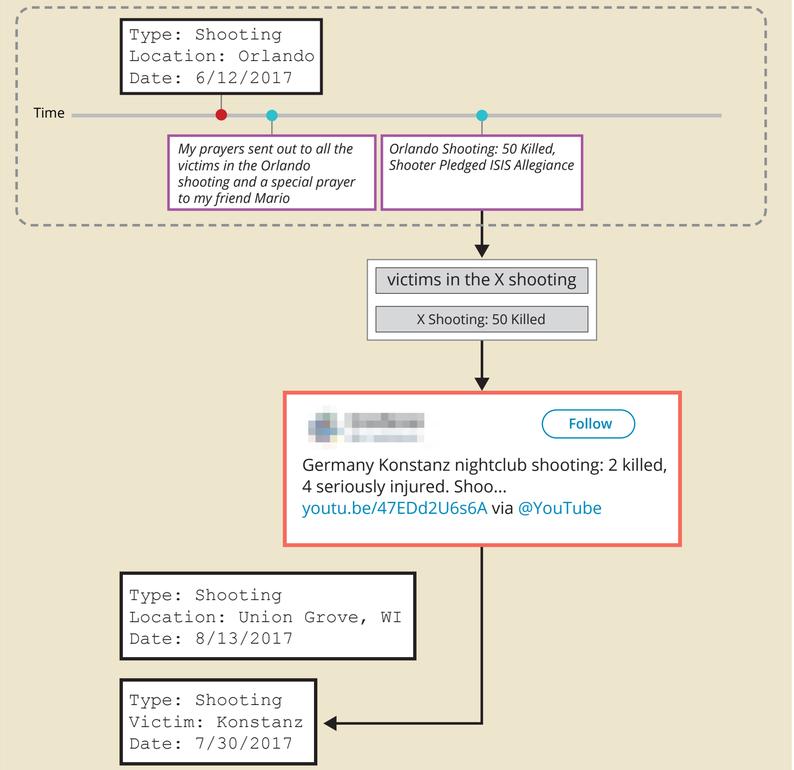

Expectation Regularization

$$O(\theta) = \underbrace{\sum_{i}^{N} \log p_\theta(y_i|x_i)}_{\text{Log Likelihood}} - \underbrace{\lambda^U D(\tilde{p}||\hat{p}_\theta^{\text{unlabeled}})}_{\text{Label regularization}}$$

Contact: Edwin Morris ejm@sei.cmu.edu

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM17-0731
Events, Relationships, and Script Learning