Sub-Event Detection from Twitter Streams as a Sequence Labeling Problem

GHENT UNIVERSITY

IDLab INTERNET & DATA LAB

G. Bekoulis, J. Deleu, T. Demeester, C. Develder

giannis.bekoulis@UGent.be

https://ugentt2k.github.io

Introduction

Goal: Detect sub-events in social media streams (i.e., Twitter)

Motivation:

(i) Difficult to track sub-events in Twitter streams

(ii) Different perspectives of the same event (e.g., emergency situations) compared to traditional media

Task: Sub-event detection in sport Twitter stream

Idea

Challenges:

(i) Noisy nature of Twitter streams (e.g., event tweets interspersed with others, non-event related info interjected)

(ii) People reporting the same thing

Contribution:

(i) Frame the problem as a sequence labeling task

(ii) Exploit the use of a chronological LSTM

Model

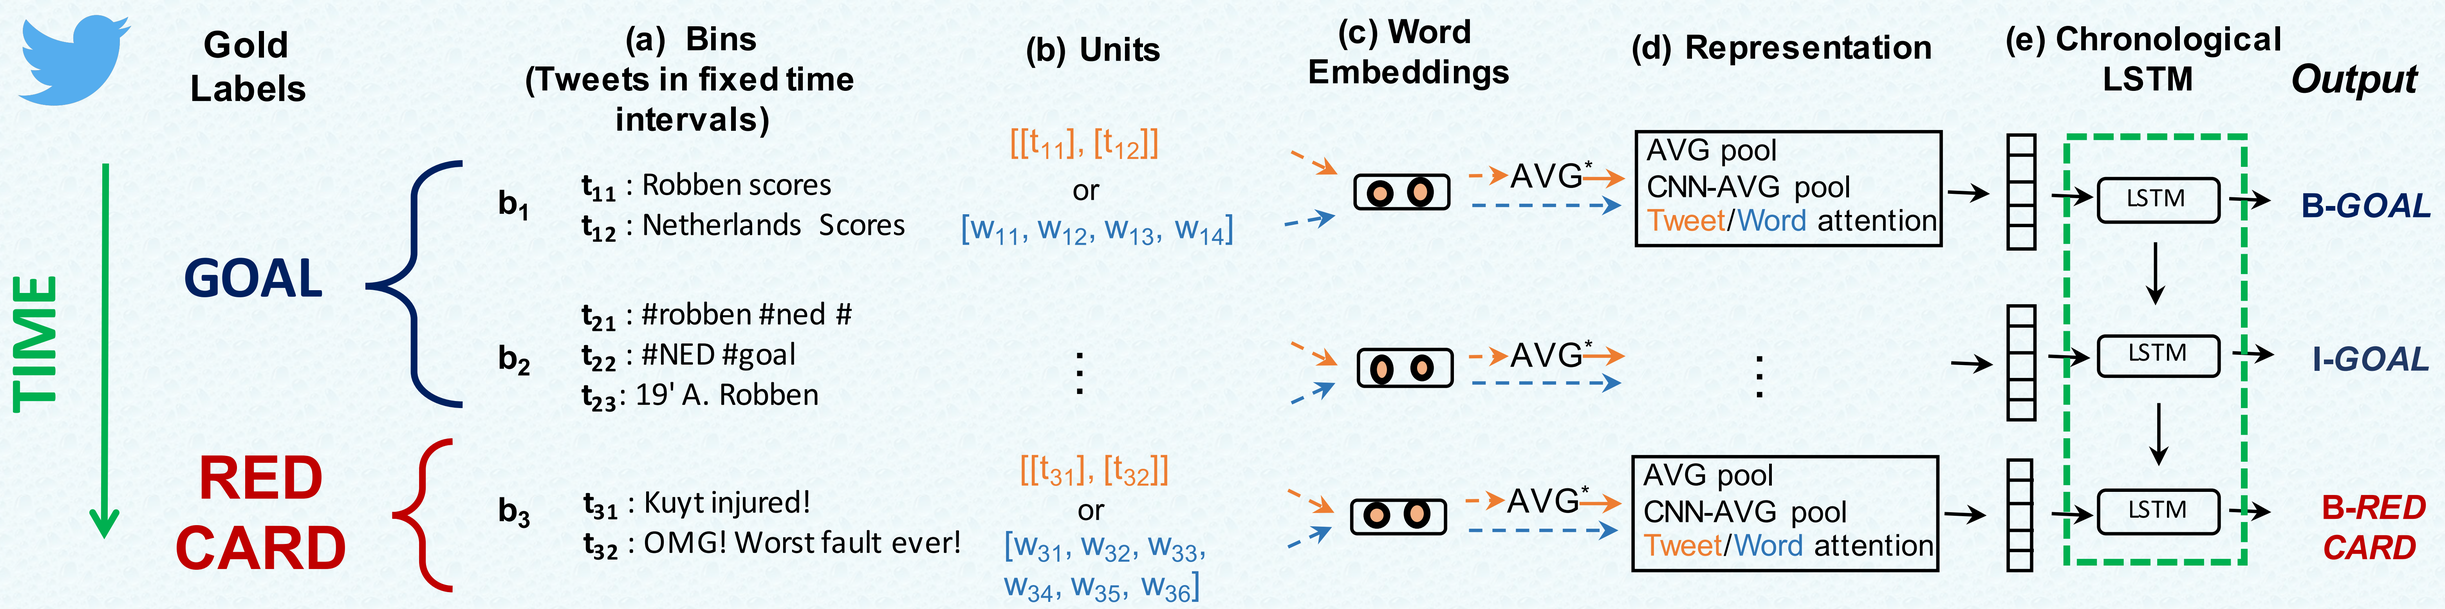

Chronological LSTM is able to capture the natural flow of the text

Results

	Model	Bin-level TL	Bin-level P	Bin-level R	Bin-level F_1	Relaxed TL	Relaxed P	Relaxed R	Relaxed F_1
without chronol. LSTM	Word-tf-idf	-	49.40	52.06	50.69	-	56.10	56.10	56.10
	Word-AVG	-	51.40	45.96	48.53	-	56.10	56.10	56.10
	Word-CNN-AVG	-	56.93	56.01	56.47	-	75.60	75.60	75.60
	Word-attention	-	52.92	58.71	55.66	-	86.59	86.59	**86.59**
	Tweet-AVG	✓	49.04	45.96	47.45	✓	62.19	62.19	62.19
	Tweet-attention	✓	51.99	42.37	46.68	✗	80.48	80.48	80.48
	Tweet-CNN	✗	58.88	51.17	54.75	✗	70.73	70.73	70.73
with chronol. LSTM	Word-AVG	-	58.14	58.35	58.24	-	71.95	71.95	71.95
	Word-CNN-AVG	-	60.89	56.19	58.45	-	60.97	60.97	60.97
	Word-attention	-	52.99	42.90	47.42	-	60.97	60.97	60.97
	Tweet-AVG	✗	57.43	60.32	**58.84**	✗	64.63	64.63	64.63
	Tweet-attention	✓	48.26	52.24	50.17	✗	67.07	67.07	67.07
	Tweet-CNN	✗	65.33	49.73	56.47	✗	60.97	60.97	60.97

Chronological LSTM improves the *bin-level* F1 score

Why is Relaxed evaluation flawed?

Gold Label	Predictions	
B-GOAL	B-RED CARD	✓
I-GOAL	I-GOAL	
I-GOAL	I-PENALTY	
I-GOAL	I-YELLOW CARD	

Rationale: If a model assigns a different label to each of the bins of a sub-event, then this sub-event counts as a true positive

Conclusions

(i) New neural model for binary sub-event detection

(ii) Propose a strong model to predict sub-event types

(iii) Extend the model with the idea of exchanging chronological information between sequential posts, and

(iv) Using a chronological LSTM is beneficial in almost all examined architectures.

References

(i) Sub-event detection from twitter streams as a sequence labeling problem. G. Bekoulis, J. Deleu, T. Demeester, C. Develder. 2019, NAACL '19.

(ii) Joint entity recognition and relation extraction as a multi-head selection problem. G. Bekoulis, J. Deleu, T. Demeester, C. Develder. 2019. Expert Systems with Applications, Volume 114, 2018.