2017-6

Tackling the Interleaving Problem in Activity Discovery

Eoin Rogers
Technological University Dublin, eoinrogers@gmail.com

Robert J. Ross
Technological University Dublin, robert.ross@tudublin.ie

John D. Kelleher
Technological University Dublin, john.d.kelleher@tudublin.ie

Recommended Citation

Rogers, E., Ross, R.J., & Kelleher, J.D. (2017). Tackling the interleaving problem in activity discovery. Trends in Cyber-Physical Multi-Agent Systems. *The PAAMS Collection - 15th International Conference*, PAAMS 2017. doi:10.1007/978-3-319-61578-3_47

Tackling the interleaving problem in activity discovery

Eoin Rogers, Robert J. Ross, and John D. Kelleher

Applied Intelligence Research Centre, Dublin Institute of Technology, Dublin, Ireland
eoin.rogers@student.dit.ie, robert.ross@dit.ie, john.d.kelleher@dit.ie

Activity discovery (AD) is the unsupervised process of discovering *activities* in data produced from streaming sensor networks that are recording the actions of human subjects. One major challenge for AD systems is *interleaving*, the tendency for people to carry out multiple activities at a time a parallel. Following on from our previous work [4], we continue to investigate AD in interleaved datasets, with a view towards progressing the state-of-the-art for AD.

A number of approaches to AD have already been published in the literature. Cook et al. [2] provide an overview of the field as it stood a number of years ago. More recently, [7] proposes a system based on clustering techniques which eliminates the need to store historic sensor data, which could help combat the data-heaviness of AD. Although not strictly a pure AD system, [6] proposes a system to take partially labelled data and automatically annotate the remainder of it. The system uses *Latent Dirichlet Allocation* (LDA), a technique from the NLP community that we also utilised for the system we mentioned above.

We are currently in the process of developing a new AD system from scratch based on deep neural networks [3]. Specifically, we generalise [1]'s concept of a *neural language model* (NLM). Given a set of n contiguous words in a sentence, a traditional NLM predicts the $n+1$th word. By contrast, our system will predict a probability distribution over the next m words. This will be used as the basis of a system that allows us to detect and remove the interleaving from a dataset automatically. This system is presented in Figure 1.

In Figure 1(a), we see the initial setup of the system: events **A** and **B** are contained in the sliding window that will constitute our input (so $n = 2$), and events **C** to **F** are in the window of length m that we want to compare to our output layer probabilities. Suppose that our language model predicts that event **D** will appear within the next m events with high probability. We thus add a *link* connecting events **B** and **D** as shown in Figure 1(b). Note that it is also possible for more links to be added. For example, if event **E** was also predicted with high probability, we would add another link from **B** to **E**. The exact mechanism by which a probability will be decided to be high enough to form a link has yet to be fully determined by the authors, although it will probably involve identifying probabilities more than a certain threshold above the average for the dataset (i.e. significantly more probable than background noise). Once all activities have been found, they can be replaced with new placeholder events, and the system can be run again (Figure 1(c)). This allows us to build a complex hierarchy of aggregate events.

Over the coming months, our intention is to implement this system and report our results back to the wider community.

References

1. Bengio, Y., Ducharme, R., Vincent, P. and Jauvin, C., 2003. A neural probabilistic language model. Journal of Machine Learning Research, 3(Feb), pp.1137-1155.
2. Cook, D.J., Krishnan, N.C. and Rashidi, P., 2013. Activity discovery and activity recognition: A new partnership. IEEE transactions on cybernetics, 43(3), pp.820-828.
3. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. Nature 521, 436–444 (2015)
4. Rogers, E., Kelleher, J.D., Ross, R.J.: Using Topic Modelling Algorithms for Hierarchical Activity Discovery. In Ambient Intelligence–Software and Applications–7th International Symposium on Ambient Intelligence (ISAmI 2016), 41–48 (2016)
5. Stoia, L., Shockley, D., Byron, D., Fosler-Lussier, E.: SCARE: A Situated Corpus with Annotated Referring Expressions. Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08) (2008)
6. Wen, J. and Wang, Z.: Learning general model for activity recognition with limited labelled data. Expert Systems with Applications, Volume 74, pp. 19–28 (2017)
7. Ye, J., Fang, L. and Dobson, S., 2016. Discovery and recognition of unknown activities. In Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (pp. 783-792). ACM.

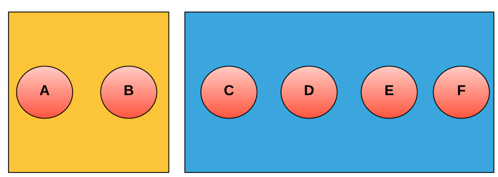

(a) The initial setup.

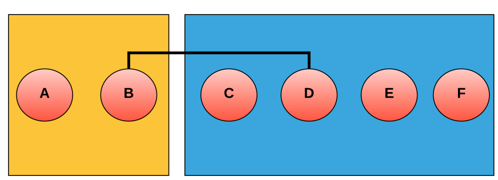

(b) A link connecting events **B** and **D**.

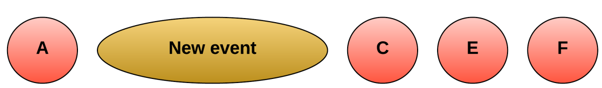

(c) The discovered activity has been subsumed into a new event.

Fig. 1. An overview of our proposed approach.