Siamese Cookie Embedding Networks for Cross-Device User Matching

Ugo Tanielian
u.tanielian@criteo.com
Criteo Research

Anne-Marie Tousch
am.tousch@criteo.com
Criteo Research

Flavian Vasile
f.vasile@criteo.com
Criteo Research

ACM Reference Format:

Ugo Tanielian, Anne-Marie Tousch, and Flavian Vasile. 2018. Siamese Cookie Embedding Networks for Cross-Device User Matching. In *WWW '18 Companion: The 2018 Web Conference Companion, April 23-27, 2018, Lyon, France*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3184558.3186941>

ABSTRACT

Over the last decade, the number of devices per person has increased substantially. This poses a challenge for cookie-based personalization applications, such as online search and advertising, as it narrows the personalization signal to a single device environment. A key task is to find which cookies belong to the same person to recover a complete cross-device user journey. Recent work on the topic has shown the benefits of using unsupervised embeddings learned on user event sequences. In this paper, we extend this approach to a supervised setting and introduce the *Siamese Cookie Embedding Network (SCEmNet)*, a siamese convolutional architecture that leverages the multi-modal aspect of sequences, and show significant improvement over the state-of-the-art.

1 INTRODUCTION

The cross-device matching task has been addressed recently in the literature, most notably within the scope of the ICDM 2015¹ challenge and the CIKM Cup 2016².

The winners of the CIKM Cup [4, 5] proposed using neural features to describe sequences of URL tokens. Their approach is based on a hierarchical *doc2vec* model [2]. However, in real world settings, *doc2vec* suffers from several drawbacks: first, *doc2vec* generates unsupervised sequence embeddings that are not specialized for the final task; secondly, *doc2vec* generates embeddings only for sequences available at training time.

In our approach, we propose to use the *TextCNN* structure as in Kim [3] in a siamese network to learn cookie similarities with different types of sequences. Using a supervised approach allows us to reach better performance than the state-of-the-art. Furthermore, our approach is able to match cookies unseen at training time.

Our main contribution is a new architecture for the problem of cookie matching, the *SCEmNet* structure, and its wide&deep counterpart (*JSCEmNet*) that address the shortcomings of previous embedded-based works.

¹ICDM 2015: <https://www.kaggle.com/c/icdm-2015-drawbridge-cross-device-connections>

²CIKM 2016: <http://cikm2016.cs.iupui.edu/cikm-cup/>

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '18 Companion, April 23-27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5640-4/18/04.

<https://doi.org/10.1145/3184558.3186941>

2 OUR APPROACH

2.1 Context

We evaluate our approach on a cross-device matching data set that was released for the CIKM Cup in 2016. The description of the data set can be found in Phan et al. [5]. The data set is composed of about 339k cookies and more than 506k ground truth pairs in training. The task is to find the 215k true pairs linking the users in test. For each cookie, the data set contains a sequence of events, where an event is a pair (URL, timestamp). The evaluation metrics of the challenge are precision, recall, and F1-score.

For a given set of cookies, the aim is then to compute a list of pairs of matched users. Since evaluating all possible pairs between n cookies results in $O(n^2)$ operations, a classical 2-step approach is to have a candidate pair generation followed by pairwise ranking trained on ground truth data.

In the first step, we use a cookie representation similar to Phan et al. [5] with k-nearest-neighbors to generate the candidate pairs. For the ranking step, we first implemented a baseline ranking algorithm. For each of the candidate pairs, we generate a set of *Baseline features*: some features are based on URL sequences such as TF, TF-IDF similarities, term-matching features and *doc2vec* embedding similarities, while other features are time-related features. For better results, we introduce our supervised *SCEmNet* architecture.

2.2 SCEmNet Architecture

Our similarity learning architecture is based on convolutional networks and is illustrated in Figure 1. We rely on the assumption that each cookie c is associated with M sequences ($s_1(c), \dots, s_M(c)$) of fixed size. For a given token modality, the sequence of tokens is first processed through a *SeqCNN* module, similar to the *TextCNN* architecture introduced in Kim [3]: given an input sequence of embedded tokens, a convolutional network with one convolution, max-pooling and dropout layer outputs a *cookie embedding*. We use this module in a siamese network with a pairwise fusion layer to transform the two cookie embeddings into a single *pair embedding*. Therefore, each one of the M modalities yields a distinct pair embedding. The M pair embeddings are combined in the multi-modal fusion layer. Finally, an output layer is learned to link the multi-modal pair embedding to the final similarity score.

The *SCEmNet* structure can therefore be trained independently to output a likelihood that a given pair is a true pair and then be used as a feature in addition to the other pre-computed *Baseline features* in a logistic regression. Or, similarly to the Wide & Deep architecture used in Cheng et al. [1], we propose to learn *SCEmNet* jointly with the logistic regression weights of the *Baseline features*: we call this the *JSCEmNet* structure.

Figure 1: SCEmNet architecture: fusionning several pair embeddings to get a richer one.

Method	Precision	Recall	F1
Baseline	40.01	42.12	41.04 $\pm$ 0.04
Baseline + SCEmNet	41.20	44.01	42.56 $\pm$ 0.05
<i>JSCEmNet</i>	44.81	47.86	46.28 $\pm$ 0.03
Baseline + SCEmNet + XGB	43.85	46.85	45.30 $\pm$ 0.05
<i>JSCEmNet</i> with XGB	44.83	47.89	46.79 $\pm$ 0.03
Phan et al. [5]	39.30	55.10	45.90

Table 1: Note that these results are computed using only 50% of the correct pairs as in the CIKM cup setup. Training the SCEmNet structure jointly with other features is the most efficient method. Here, Baseline refers to the Baseline features defined previously.

3 EXPERIMENTS AND RESULTS

For the *SCEmNet* implementation, we fixed the different filter sizes to be (2, 3, 4, 5, 6, 10) and the number of filters per size at 40. The dropout probability was set at 0.5 for training and 0 for test. Negative pairs were sampled from the set of candidate pairs not present in the list of ground truth pairs. Training a multi-modal *SCEmNet*, where $M = 4$ with sequences generated from URLs at different depths, resulted in learning richer pair embeddings. Regarding the merging of cookie embeddings, we found that using element wise multiplication for the fusion layer and concatenation for the multi-modal layer improved the results.

We compare the performance of the *SCEmNet* feature with *JSCEmNet* using only 50% of the test pairs, following the CIKM challenge setup. We report confidence intervals on the F1 score obtained by doing 50 random splits of the test pairs. Note that this is different to how Phan et al. [5] reported their results, as they evaluated on the full test set.

Our main results are presented in Table 1. We show that using all features in joint training is the most efficient method: *JSCEmNet* is doing better than the baseline features combined with *SCEmNet*. Our best score was obtained when adding a gradient boosted feature (XGB), which was obtained by training an XGBoost model on the baseline features and using its prediction as a feature. We improved the state-of-the-art results of Phan et al. [5] by 1.9%.

4 CONCLUSION

This paper adds to the literature showing that applying Natural Language Processing techniques to sequences can be very efficient, especially in the field of similarity learning. We introduced the *SCEmNet* structure showing that siamese CNNs can be applied successfully to learn sequence similarity on non textual data. Furthermore, with the *JSCEmNet* architecture, we showed that learning the *SCEmNet* representation jointly with other features brings additional value and beats state-of-the-art results.

These promising results were obtained while keeping to state-of-the-art methods for candidate generation. Besides, when working on online users, training a sequence representation such as *doc2vec* for all cookies might not be realistic (e.g. billions of sequences). In that case it could be helpful to use a *SCEmNet*-like structure to infer sequence embeddings and perform approximate kNN search. Future work should also better integrate the temporal information (time gaps, daily and weekly patterns,...) in our cookie/pair representation. A possibility could be to integrate the time-gaps in *doc2vec* as additional information as in Vasile et al. [7], or to use RNNs as proposed by Smirnova and Vasile [6]. Further work should investigate the use of RNNs for the task of similarity learning.

More generally, this structure developed for cookie matching can be generalized to any kind of sequence similarity task. For example, we plan to test our similarity learning structure on textual data for the SemEval-2018 workshop on Semantic Evaluation³.

REFERENCES

- [1] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishii Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ipsir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*. ACM, 7–10.
- [2] Andrew M. Dai, Christopher Olah, and Quoc V. Le. 2015. Document Embedding with Paragraph Vectors. *CoRR* abs/1507.07998 (2015). <http://arxiv.org/abs/1507.07998>
- [3] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. *arXiv preprint arXiv:1408.5882* (2014). arXiv:arXiv:1408.5882v2
- [4] Cong-Minh Phan, Yin Tay, and Tuan-Anh Nguyen Pham. 2016. Cross Device Matching for Online Advertising with Neural Feature Ensembles : First Place Solution at CIKM Cup 2016. In *CIKM 2016*.
- [5] Min h C. Phan, Aixin Sun, and Yi Tay. 2017. Cross-Device User Linking : URL, Session, Visiting Time, and Device-log Embedding. In *SIGIR*. 933–936.
- [6] Elena Smirnova and Flavian Vasile. 2017. Contextual Sequence Modeling for Recommendation with Recurrent Neural Networks. In *WDLRS*. arXiv:1706.07684v1
- [7] Flavian Vasile, Elena Smirnova, and Alexis Conneau. 2016. Meta-Prod2Vec - Product Embeddings Using Side-Information for Recommendation. *RecSys* (2016), 0–7. arXiv:1607.07326

³SemEval-2018: <http://alt.qcri.org/semeval2018/>