

# Hierarchical Graph Neural Network with Cross-Attention for Cross-Device User Matching

Ali Taghibakhshi<sup>1,2,\*</sup>, Mingyuan Ma<sup>1</sup>, Ashwath Aithal<sup>1</sup>, Onur Yilmaz<sup>1</sup>, Haggai Maron<sup>1</sup>, and Matthew West<sup>2</sup>

<sup>1</sup> NVIDIA

<sup>2</sup> Department of Mechanical Science and Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA

**Abstract.** Cross-device user matching is a critical problem in numerous domains, including advertising, recommender systems, and cybersecurity. It involves identifying and linking different devices belonging to the same person, utilizing sequence logs. Previous data mining techniques have struggled to address the long-range dependencies and higher-order connections between the logs. Recently, researchers have modeled this problem as a graph problem and proposed a two-tier graph contextual embedding (TGCE) neural network architecture, which outperforms previous methods. In this paper, we propose a novel hierarchical graph neural network architecture (HGNN), which has a more computationally efficient second level design than TGCE. Furthermore, we introduce a cross-attention (Cross-Att) mechanism in our model, which improves performance by 5% compared to the state-of-the-art TGCE method.

**Keywords:** Graph neural network · User matching · Cross-attention.

## 1 Introduction

Ensuring system security and effective data management are critical challenges in the modern day [3, 4]. In this regard, data integration plays a vital role in facilitating data management, as it enables the integration of data from diverse sources to generate a unified view of the underlying domain. One of the primary challenges in data integration is the problem of entity resolution, which involves identifying and linking multiple data records that correspond to the same real-world entity. The problem of entity resolution arises in a wide range of domains, including healthcare, finance, social media, and e-commerce. Entity resolution is a challenging problem due to various factors, including the presence of noisy and ambiguous data, the lack of unique identifiers for entities, and the complexity of the relationships between different entities.

Among entity resolution tasks, cross-device user matching is of significant importance. This task involves determining whether two separate devices belong to the same real-world person based on their sequential logs. The device

---

\*This work was done while Ali Taghibakhshi was an intern at NVIDIA.

sequential logs are time-stamped actions taken by the user over a relatively long period of time, say a few months. These actions are often in the form of browsing a Uniform Resource Locator (URL), and almost always, user identifications are not available due to privacy reasons. Refer to Figure 1 for an illustration of the cross-device user matching task.




Fig. 1: Cross device user matching problem: only based the URL visit logs of two different devices, determine whether or not they belong to the same real-world person.

It is a common occurrence for users to engage in online activities across multiple devices. However, businesses and brands often struggle with having insufficient user identities to work with since users are perceived as different individuals across different devices due to their unique activities. The ability to automatically identify the same user across multiple devices is essential for gaining insights into human behavior patterns, which can aid in applications such as user profiling, online advertising, improving system security. Therefore, in recent years, there has been a flourishing amount of studies focusing on cross-device user matching [9].

In recent years, with the advent of machine learning-based methods for entity resolution, several studies have focused on learning distributed embeddings for the devices based on their URL logs [6, 11, 12]. The earlier studies focused on utilizing unsupervised feature learning techniques [7], developing handcrafted features for the device logs, or relied on co-occurrence of key attributes of URL logs in pairwise classification [12].

Methods that utilize deep learning have a greater ability to convey dense connections among the sequential device logs. For instance, researchers have utilized a 2D convolutional neural network (CNN) framework to encode sequential log representations to understand the relationship between two devices [16]. However, this model primarily captures local interactions within user sequence logs, limiting its ability to learn the entire sequence or a higher-level pattern. Recently, there has been further emphasis on the effectiveness of sequential models like recurrent neural networks (RNNs) and attention-based techniques in modeling sequence patterns and achieving promising results in numerous sequence modeling tasks [5, 13, 15]. Although these methods work well for sequence modeling, they are not specifically designed for user-matching tasks and may not be optimal for learning sequential log embeddings.

Recently, researchers proposed a two-tier graph contextual embedding (TGCE) network for the cross-device user matching [6] task. While previous methods for the task often failed at long-range information passing along the sequence logs, TGCE leverages a two-level structure that can facilitate information passing beyond the immediate neighborhood of a device log. This was specifically achieved by considering a random walk starting from every node in a device log, connecting all of the visited nodes to the original node, and performing a round of message passing using the newly generated shortcut edges.

Although the two-tier structure seems to enable long-range information sharing, we note two major limitations with the existing method. First, in the device graph, the random walk on the URL nodes may randomly connect two URLs that have been visited at two far-away time-stamps. Intuitively, two different URLs browsed by a device with weeks of gap in between share less information than two URLs visited in a shorter time frame. Second, at the end of the TGCE architecture, for the pairwise classification task, the generated graph embeddings for two devices are entry-wise multiplied and sent through a fully connected network to determine if they belong to the same person. However, there could be significant key features in the learned embeddings that may be shared between the devices, which can alternatively get lost if the architecture does not compare them across one another.

To address the above two issues, we propose a new hierarchical graph neural network (HGNN) inspired by the star graph architecture [10]. In the terminology of HGNN, we refer to the URL nodes as *fine* nodes, and in an unraveled sequence of URL logs, HGNN assigns a *coarse* node to every  $K$  consecutive fine nodes. The message passing between the coarse and fine nodes enables effective long-range message passing without the need to excessively add edges, as in the random walk method. Moreover, for the pairwise classification task, we utilize a cross-attention mechanism inspired by Li *et. al.* [8], which enables entry-wise cross-encoding of the learned embeddings. The main contributions of this paper are summarized as follows:

- We model a given device log as a hierarchical heterogeneous graph, which is 6x faster than the previous state-of-the-art while keeping a competitive level of accuracy and performance.
- We employ a cross-attention mechanism for pairwise matching of the graphs associated with a device log, which improves the accuracy of the overall method by about 5%.

## 2 Related work

The cross-device user matching task was first introduced in the CIKM Cup 2016<sup>‡</sup>, and the first proposed methods for the task mainly considered hand-crafted features. For instance, the runner-up solution [9] produces sub-categories based on the most significant URLs to generate detailed features. Furthermore, the

<sup>‡</sup><http://cikm2016.cs.iupui.edu/cikm-cup/>

competition winner solution proposed by Tay *et. al.* [14] utilizes “term frequency inverse document frequency” (TD-IDF) features of URLs and other related URL visit time features. However, their manually designed features did not fully investigate more intricate semantic details, such as the order of behavior sequences, which restricted their effectiveness. Aside from the hand-crafted features, the features that are developed from the structural information of the device URL visit data are also crucial for accomplishing the task of user matching. To further process sequential log information, studies have applied LSTM, 2D-CNN, and Doc2vec to generate semantic features for a sequence visited by a device [11, 12, 16].

Sequence-based machine learning models have also been employed for different entity resolution tasks; for instance, recurrent neural networks (RNN) have been utilized to encode behavior item sequential information [5]. Nevertheless, long-range dependencies and more advanced sequence features are not well obtained using sequence models [6]. With the advent of graph neural networks (GNN), researchers have modeled device logs as individual graphs where nodes and edges represent visited URLs and transitions between URLs. Each node and/or edge has an initial feature vector obtained from the underlying problem, and the layers of GNNs are then employed to update these features based on information passing in the local neighborhood of every node, such as the SR-GNN paper [17]. Another example is the LESSR [1] method for recommendation systems where the method is capable of long-range information capturing using an edge-order preserving architecture. However, these methods are specifically designed for the recommendation task and do not necessarily achieve desirable results on the cross-device user matching task.

Recently, researchers have proposed TGCE [6], a two-tier GNN for the cross-device user matching task. In the first tier, for every device log, each URL is considered as a node, and directional edges denote transitions between URLs. In the second tier, shortcut edges are formed by starting a random walk from every node and connecting all of the visited nodes to it. After a round of message passing in the first tier, the second tier is supposed to facilitate long-range information sharing in the device log. After the second tier, a position-aware graph attention layer is applied, followed by an attention pooling, which outputs the learned embedding for the whole graph. For the final pairwise classification, these learned embeddings for each of the devices are multiplied in an entry-wise manner and are sent to a fully connected deep neural network to determine whether they belong to the same user.

### 3 Hierarchical graph neural network

In this section, we discuss how we employ a two-level heterogeneous graph neural network for the cross-device user matching problem.

### 3.1 Problem definition

The aim of the cross-device user matching problem is to determine whether two devices belong to the same user, given only the URL visits of each device. Denote a sequence of visited URLs by a device  $v$  by  $\mathcal{S}_v = \{s_1, s_2, \dots, s_n\}$ , where  $s_i$  denotes the  $i$ 'th URL visit by the device (note that  $s_i$ 's are not necessarily different). We build a hierarchical heterogeneous graph,  $G_v$ , based on the sequence  $\mathcal{S}_v$  as follows: for a visited URL,  $s_i$ , consider a *fine* node in  $G_v$  and denote it by  $f_i$ . Note that if multiple  $s_i$ 's correspond to the same URL, we only consider one node for it in  $G_v$ . Then, we connect nodes corresponding to consecutively visited URLs by directed edges in the graph; we connect  $f_i$  and  $f_{i+1}$  by a directional edge (if  $f_i$  and  $f_{i+1}$  correspond to the same URL, the edge becomes a self-loop). Up to this point, we have defined the fine-level graph, and we are ready to construct the second level, which we call the *coarse* level.

To construct the second level, we partition the sequence  $\mathcal{S}_v$  into non-overlapping subgroups of  $K$  URLs, where each subgroup consists of consecutively visited URLs (the last subgroup may have less than  $K$  URLs). For every subgroup  $j$ , we consider a coarse node,  $c_j$ , and connect it to all of the fine nodes corresponding to the URLs in subgroup  $j$  via undirected edges.

### 3.2 Fine level

In the fine level of the graph  $G_v$ , for every node  $f_i$ , we order the nodes corresponding to the URLs that have an incoming edge to  $f_i$  according to their position in  $\mathcal{S}_v$ . We denote this ordered sequence of nodes by  $N_i = \{f_{j_1}, f_{j_2}, \dots, f_{j_\kappa}\}$ . Also, we denote the feature vector of the fine node  $f_i$  by  $x_i$ . The  $l$ -th round of message passing in the fine-level graph updates the node features according to the following update methods:

$$M_i^{(l)} = \Phi^{(l)}([x_{j_1}, x_{j_2}, \dots, x_{j_\kappa}, x_i]), \quad (1)$$

$$x_i^{(l+1)} = \Psi^{(l)}(x_i^{(l)}, M_i^{(l)}), \quad (2)$$

where  $\Phi^{(l)}$  is a sequence aggregation function (such as sum, max, GRU, LSTM, etc.), for which we use GRU [2], and  $\Psi^{(l)}$  is a function for updating the feature vector (e.g., a neural network), for which we use a simple mean.

### 3.3 Coarse level

In every round of heterogeneous message passing between fine and coarse level nodes, we update both the fine and coarse node features. Consider the coarse node  $c_j$ , and denote its feature by  $\tilde{x}_j$ . Also, denote the fine neighbor nodes of  $c_j$  by  $\tilde{\mathcal{N}}(c_j)$ . In the  $l$ -th layer of heterogeneous message passing, the coarse node feature update is as follows:

$$\tilde{x}_j^{(l+1)} = \square_{i \in \tilde{\mathcal{N}}(c_j)} (W_1^{(i)} x_i), \quad (3)$$

where  $W_1^{(l)}$  is a learnable matrix and  $\square$  is an aggregation function (such as mean, max, sum, etc.), for which we use mean. Denote by  $\mathcal{N}(f_i)$  the set of coarse nodes connected to the fine node  $f_i$ . We first learn attention weights for the heterogeneous edges, and then we update fine nodes accordingly. In the  $l$ -th round of heterogeneous message passing, the fine node features are updated as follows:

$$e_{i,j}^{(l)} = \phi(W_2^{(l)} x_i^{(l)}, W_3^{(l)} \tilde{x}_j^{(l)}), \quad (4)$$

$$\alpha_{i,j}^{(l)} = \frac{\exp(e_{i,j}^{(l)})}{\sum_{j \in \mathcal{N}(f_i)} \exp(e_{i,j}^{(l)})}, \quad (5)$$

$$x_i^{(l+1)} = \xi(x_i^{(l)}, \sum_{j \in \mathcal{N}(f_i)} \alpha_{i,j}^{(l)} \tilde{x}_j^{(l)}), \quad (6)$$

where  $W_2^{(l)}$  and  $W_3^{(l)}$  are learnable matrices, and  $\xi$  and  $\phi$  are update functions (such as a fully connected network). Figure 2 shows the overall architecture of fine and coarse level message passing.




Fig. 2: From left to right: heterogeneous (fine and coarse) graph modeling from a given URL sequence. The hierarchical message passing blocks consist of message passing on the fine nodes with a GRU aggregation function. Next, the coarse node features are updated using a mean aggregation function. Finally, the fine node features are updated using their previous feature vector as well as an aggregated message from their associated coarse nodes obtained via an attention mechanism between coarse and fine level nodes.

### 3.4 Cross attention

After the message passing rounds in the fine level and long-range information sharing between fine and coarse nodes, we extract the learned fine node embeddings and proceed to cross encoding and feature filtering, inspired by the

GraphER architecture [8]. We consider two different device logs  $v$  and  $w$ , and their learned fine node embeddings as a sequence, ignoring the underlying graph structure. We denote the learned fine node embeddings for device logs  $v$  and  $w$  as  $X_v \in \mathbb{R}^{m_v \times d}$  and  $X_w \in \mathbb{R}^{m_w \times d}$ , where  $m_v$  and  $m_w$  are the number of nodes in the fine level of  $G_v$  and  $G_w$ , respectively. We learn two matrices for cross-encoding  $X_v$  into  $X_w$  and vice-versa. Consider the  $i$ -th and  $j$ -th rows of  $X_v$  and  $X_w$ , respectively, and denote them by  $x_{v,i}$  and  $x_{w,j}$ . The entries  $\hat{\alpha}_{i,j}$  of the matrix  $A_{v,w}$  for cross-encoding  $X_v$  into  $X_w$  are obtained using an attention mechanism (and similarly for  $A_{w,v}$ ):

$$\hat{e}_{i,j} = \zeta(W_3 x_{v,i}, W_3 x_{w,j}), \quad (7)$$

$$\hat{\alpha}_{i,j} = \frac{\exp(\hat{e}_{i,j})}{\sum_{k=1}^{m_w} \exp(\hat{e}_{i,k})}, \quad (8)$$

where  $\zeta$  is an update function (such as a neural network), for which we use a simple mean. After obtaining the cross-encoding weights, we apply feature filtering, a self-attention mechanism that filters important features. The filtering vector is obtained as  $\beta_v = \text{sigmoid}(W_4 \tanh(W_5 X_v^T))$ , where  $W_4$  and  $W_5$  are learnable weights ( $\beta_w$  is obtained similarly). We apply the feature-filtering vector to the cross-encoding matrix as follows:

$$L_{v,w} = [\text{diag}(\beta_v)(A_{v,w}X_w - X_v)] \odot [\text{diag}(\beta_v)(A_{v,w}X_w - X_v)], \quad (9)$$

where  $\odot$  denotes the Hadamard product ( $L_{w,v}$  is also obtained similarly). The  $L_{v,w} \in \mathbb{R}^{m_v \times d}$  and  $L_{w,v} \in \mathbb{R}^{m_w \times d}$  matrices come from the Euclidean distance between the cross-encoding of  $X_v$  into  $X_w$  and  $X_w$ , and therefore are a measure of the closeness of the original sequence logs of  $v$  and  $w$ .

To obtain a size-independent comparison metric, we apply a multi-layer perceptron (MLP) along the feature dimension of  $L$  matrices (the second dimension,  $d$ ), followed by a max-pooling operation along the first dimension. Finally, we apply a dropout and a ReLU nonlinearity. This yields vectors  $r_{v,w}$  and  $r_{w,v}$  that have a fixed size for any pair of  $v$  and  $w$ . For the final pairwise classification task, we concatenate  $r_{v,w}$  and  $r_{w,v}$  and pass it through an MLP followed by a sigmoid activation to determine if the two devices belong to the same user or not:

$$\hat{y} = \text{sigmoid}(\text{MLP}(r_{v,w} || r_{w,v})). \quad (10)$$

## 4 Experiment

In this section, we will describe the dataset, training details, and discuss how our method outperforms all other baselines, including TGCE [6], the previous state-of-the-art.

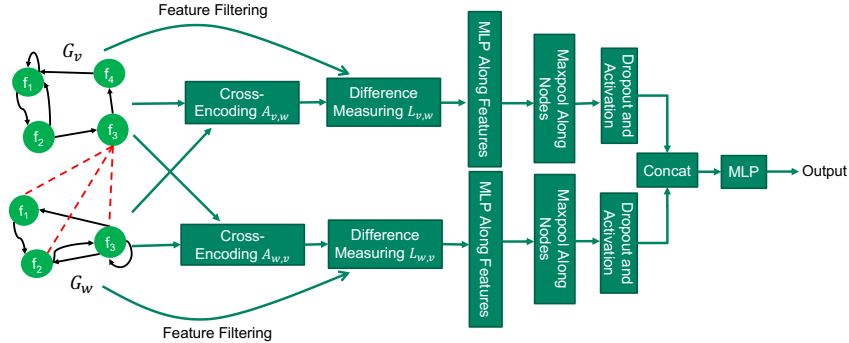


Fig. 3: Pairwise device graph matching: After the message passing, the two device graphs are cross-encoded via an attention mechanism followed by an attention-based feature filtering. The resulting matrix for each graph is then passed through an MLP layer, acting along the feature, followed by a maxpool operator along the nodes. Next, the obtained vectors pass through a dropout layer followed by an activation function. Finally, the resulting vectors of the two graphs are concatenated and passed through an MLP to obtain the final output.

#### 4.1 Training details

We studied the cross-device user matching dataset made publicly available by the Data Centric Alliance<sup>§</sup> for the CIKM Cup 2016 competition. The dataset consists of 14,148,535 anonymized URL logs of different devices with an average of 197 logs per device. The dataset is split into 50,146 and 48,122 training and test device logs, respectively. To obtain the initial embeddings of each URL, we applied the same data preprocessing methods as in [6, 11]. We used a coarse-to-fine node ratio of  $k = 6$ , a batch size of 800 pairs of device logs, a learning rate of  $10^{-3}$ , and trained the model for 20 epochs. We used the binary cross-entropy (BCE) loss function for training our model. The training, evaluation, and test were all executed on an A100 NVIDIA GPU. The BCE loss during training as well as the validation F1 score are shown in Figures 4 and 5, respectively.

#### 4.2 Results

In this section, we evaluate the precision, recall, and F1 score of our method on the test set and compare it to available baselines. All of the baselines have been obtained similarly as described in [6]. We present two variants of our method; the first one, which we label “HGNN”, only differs from TGCE in the design of the second tier, i.e., we use the hierarchical structure presented in subsections 3.2 and 3.3, followed by the rest of the TGCE architecture. The second variant, which we label “HGNN+Cross-Att”, uses the hierarchical structure in

<sup>§</sup><https://competitions.codalab.org/competitions/11171>






Fig. 4: Binary cross-entropy loss of our proposed method against that of TGCE. During training, our method obtains strictly better loss values.




Fig. 5: Validation F1 score during training. Throughout the training, our method achieves strictly better F1 scores for the validation set compared to that of TGCE.

subsections 3.2 and 3.3, and also utilizes the cross-attention mechanism presented in subsection 3.4 after the hierarchical structure. As shown in Table 1, the ‘‘HGNN+Cross-Att’’ variant outperforms all of the baselines on the F1 score metric, including the second-best method (TGCE) by 5% on the test data.

Table 1: Precision, Recall, and F1 score of different methods for cross-device user matching on DCA dataset.

	Precision at Best F1 Score	Recall at Best F1 Score	Best F1 Score
TF-IDF	0.33	0.27	0.26
Doc2vec	0.29	0.21	0.24
SCEmNet	0.38	0.44	0.41
GRU	0.37	<b>0.49</b>	0.42
Transformer	0.39	0.47	0.43
SR-GNN	0.35	0.34	0.34
LESSER	0.41	<b>0.48</b>	0.44
TGCE	0.49	0.44	0.46
HGNN (ours)	0.48	0.43	0.45
HGNN+Cross-Att (ours)	<b>0.57</b>	<b>0.48</b>	<b>0.51</b>

We also compare the training time of the two variants of our method with that of TGCE. As shown in Table 2, our hierarchical structure is significantly more efficient than that of TGCE while keeping a competitive F1 score. Table 2 essentially indicates that by simply replacing the second-tier design of TGCE with our hierarchical structure (presented in subsections 3.2 and 3.3), the method becomes 6x faster while almost keeping the same performance. This is due to the large number of artificial edges generated in the random walk passes in the creation of the second tier of TGCE. Moreover, although including cross-attention slows down the model, we can still obtain the same training time as TGCE and achieve 5% better overall F1 score.

Figure 6 shows the precision-recall curve of our method (the HGNN+Cross-Att variant, trained for 6 epochs) with that of TGCE (trained for 20 epochs). As shown in the figure, the precision-recall curve of our method is strictly better than that of TGCE. In other words, for every recall score, our method has a better precision. Additionally, we further trained the HGNN+Cross-Att variant for 20 epochs (the same number of epochs TGCE was trained for) to study if any further improvement is achieved on the test set. We also plot the F1 score with different thresholds (from 0 to 1 incremented by 0.01) for our model trained for 6 and 20 epochs and compare it to that of TGCE. As shown in Figure 7, our model trained for 20 epochs strictly outperforms TGCE (also trained for 20 epochs) for every threshold for obtaining the F1 score. However, our model

Table 2: Best F1 score and end-to-end training time of HGNN (without Cross-Att), HGNN+Cross-Att, and TGCE. The HGNN model is 6x faster than TGCE with a slight trade-off (about 1%) on the accuracy side. The HGNN+Cross-Att model has the same training time as TGCE while achieving 5% better F1 score.

	<b>Best F1 Score</b>	<b>End-to-end Training Time</b>	<b>Number of Epochs</b>
TGCE	0.46	60h	20
HGNN (ours)	0.45	<b>10h</b>	20
HGNN+Cross-Att (ours)	<b>0.51</b>	60h	<b>6</b>

trained for 6 epochs achieves the best overall F1 score, which is 5% higher than TGCE. This is significant since as shown in Table 2, the model takes the same time as TGCE to train.

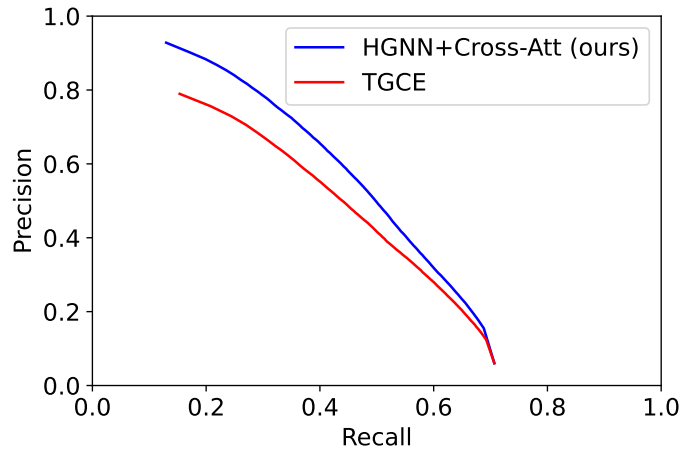


Fig. 6: Precision-Recall curve of the proposed method and that of TGCE on the test data.




Fig. 7: F1 score against rounding threshold for our method (both for networks trained for 6 and 20 epochs) compared to that of TGCE (trained for 20 epochs). When our model is trained for 6 epochs, it achieves the highest overall F1 score, which is about 5% better than that of TGCE. However, if our model is trained for 20 epochs, it achieves higher F1 scores for a wider range of rounding thresholds, outperforming TGCE for every rounding threshold.

## 5 Conclusions

In this paper, we present a novel graph neural network (GNN) architecture for a demanding entity resolution task: cross-device user matching, which determines if two devices belong to the same user based only on their anonymized internet logs. Our method comprises of designing an effective hierarchical structure for achieving long-range message passing in the graph obtained from device URL logs. After passing device logs through such a hierarchical GNN, we employ a cross-attention mechanism to effectively compare device logs against each other to determine if they belong to the same user. We demonstrate that our method outperforms available baselines by at least 5%, while having the same training time as the previous state-of-the-art method, establishing the effectiveness of our proposed method.

**Acknowledgements** This research was supported by NVIDIA Corporation.

## References

1. Chen, T., Wong, R.C.W.: Handling information loss of graph neural networks for session-based recommendation. In: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 1172–1180 (2020)

2. Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv preprint arXiv:1412.3555 (2014)
3. Gharaibeh, A., Salahuddin, M.A., Hussini, S.J., Khreishah, A., Khalil, I., Guizani, M., Al-Fuqaha, A.: Smart cities: A survey on data management, security, and enabling technologies. *IEEE Communications Surveys & Tutorials* **19**(4), 2456–2501 (2017)
4. Gholizadeh, N.: Iec 61850 standard and its capabilities in protection systems (2016)
5. Hidasi, B., Karatzoglou, A., Baltrunas, L., Tikk, D.: Session-based recommendations with recurrent neural networks. arXiv preprint arXiv:1511.06939 (2015)
6. Huang, H., Guo, S., Li, C., Sheng, J., Wang, L., Li, J., Liu, J., Zhong, S.: Two-tier graph contextual embedding for cross-device user matching. In: *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. pp. 730–739 (2021)
7. Le, Q., Mikolov, T.: Distributed representations of sentences and documents. In: *International conference on machine learning*. pp. 1188–1196. PMLR (2014)
8. Li, B., Wang, W., Sun, Y., Zhang, L., Ali, M.A., Wang, Y.: Grapher: token-centric entity resolution with graph convolutional neural networks. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 34, pp. 8172–8179 (2020)
9. Lian, J., Xie, X.: Cross-device user matching based on massive browse logs: The runner-up solution for the 2016 cism cup. arXiv preprint arXiv:1610.03928 (2016)
10. Pan, Z., Cai, F., Chen, W., Chen, H., De Rijke, M.: Star graph neural networks for session-based recommendation. In: *Proceedings of the 29th ACM international conference on information & knowledge management*. pp. 1195–1204 (2020)
11. Phan, M.C., Sun, A., Tay, Y.: Cross-device user linking: url, session, visiting time, and device-log embedding. In: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. pp. 933–936 (2017)
12. Phan, M.C., Tay, Y., Pham, T.A.N.: Cross device matching for online advertising with neural feature ensembles: First place solution at cism cup 2016. arXiv preprint arXiv:1610.07119 (2016)
13. Qiu, R., Yin, H., Huang, Z., Chen, T.: Gag: Global attributed graph neural network for streaming session-based recommendation. In: *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*. pp. 669–678 (2020)
14. Ramos, J., et al.: Using tf-idf to determine word relevance in document queries. In: *Proceedings of the first instructional conference on machine learning*. vol. 242, pp. 29–48. Citeseer (2003)
15. Sun, F., Liu, J., Wu, J., Pei, C., Lin, X., Ou, W., Jiang, P.: Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In: *Proceedings of the 28th ACM international conference on information and knowledge management*. pp. 1441–1450 (2019)
16. Tanielian, U., Tusch, A.M., Vasile, F.: Siamese cookie embedding networks for cross-device user matching. In: *Companion Proceedings of the The Web Conference 2018*. pp. 85–86 (2018)
17. Wu, S., Tang, Y., Zhu, Y., Wang, L., Xie, X., Tan, T.: Session-based recommendation with graph neural networks. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 33, pp. 346–353 (2019)