
On Embeddings in Relational Databases

Siddhant Arora Srikanta Bedathur

CSE, IIT Delhi,

Delhi, India.

{Siddhant.Arora.cs515, srikanta}@cse.iitd.ac.in

Abstract

We address the problem of learning a distributed representation of entities in a relational database using a low-dimensional embedding. Low-dimensional embeddings aim to encapsulate a concise vector representation for an underlying dataset with minimum loss of information. Embeddings across entities in a relational database have been less explored due to the intricate data relations and representation complexity involved. Relational databases are an inter-weaved collection of relations that not only model relationships between entities but also record complex domain-specific quantitative and temporal attributes of data defining complex relationships among entities. Recent methods for learning an embedding constitute of a naive approach to consider complete *denormalization* of the database by materializing the full join of all tables and representing as a knowledge graph. This popular approach has certain limitations as it fails to capture the inter-row relationships and additional semantics encoded in the relational databases. In this paper we demonstrate; a better methodology for learning representations by exploiting the underlying semantics of columns in a table while using the relation joins and the latent inter-row relationships. Empirical results over a real-world database with evaluations on similarity join and table completion tasks support our proposition.

1 Introduction

Embedding based methods have recently gathered a lot of attention mainly due to their ability to generate a mapping of categorical/discrete variables to continuous vector spaces and thus their support towards neural network models. They have been applied to a variety of tasks namely, images (Krizhevsky et al., 2012), text (Mikolov et al., 2013), knowledge graphs (Lin et al., 2015), time series (Chung et al., 2016) with limited work on effective embeddings for entities in databases with *multiple* relational tables. Bordawekar and Shmueli (2017) introduced relational databases embeddings, but their proposed approach treated the row of a table as a sentence for learning word2vec embeddings of an entity and did not consider the problem involved in handling heterogeneous relations that exist between entities or inter-row relationships encoded in the databases.

Operational databases consist of multiple tables with each table incorporating complex relations both within and across rows of tables. This complexity persists even when we consider a fully *denormalized* relational database, obtained by completing all the natural joins between primary-key–foreign-key columns. For instance, the real-world IMDB movie database (imd) contains among others, Director (id, first_name, last_name) and Directors_Genre (director_id, genre, prob) tables which maintain information with each entity consisting of a *director* with directed a movie and their genre preferences respectively. Specifically, the Directors_Genre.prob is the observed probability of the corresponding director directing a movie of that genre. Naturally, this continuous-valued attributes cannot be captured accurately by a naive word2vec model operating directly on the table. Instead, by using the semantic knowledge of the data being modeled, we can

generate a modified representation of the table to significantly boost the performance of the word2vec model.

Another possible view is to consider the table as a knowledge graph with a KG embedding method while learning representations. Apart from the limitations above, constructing a KG fails to capture inter-row relationships that are often implicit in the database and play an important role in the effective representation of entities. For instance, consider two tables `Movies` (`id`, `name`, `year`, `rank`) and `Movies_Directors` (`director_id`, `movie_id`) connecting movie details with corresponding director(s). Now, two directors with similar trajectories in their directorial outcomes are expected to close in their representations. But neither of the above two models of learning representations of relational tables will be able to capture this effect. Instead, a specialized sequence embedding model such as an RNN with LSTM (Hochreiter and Schmidhuber, 1997) captures the temporal characteristics using the inter-row sequential relationships encoded in the database.

We further support our observations through empirically designing and evaluating alternate embedding strategies that utilize the semantics of columns in a relational table. Specifically, we consider using weighted sample views of the database table (which can be implemented in SQL efficiently) as the input to word2vec models for two strategies in the context of our movie database –viz., *Director-Genre* sampling and *Movie-rank* sampling. Later, we consider LSTM for capturing inter-record relations because of the temporal data encoded in the movie database. Naturally, other alternatives may need to be used in other forms of inter-record relationships are to be captured (e.g., spatial proximity of locations).

We consider our evaluations across two metrics; (i) Similarity queries, and (ii) data cell completion and demonstrate a novel coherent evaluation framework for similarity query task by using Milne-Witten (Witten and Milne, 2008) inlink score on Wikipedia as the baseline.

To summarise our contributions in this work:

1. We overcome the limitations associated with the previous methods and leverage semantic level information to effectively learn embeddings on databases.
2. Exploiting inter-row relation and time-series data that are highly common in database settings along with a rigorous evaluation across multiple baselines.

2 Related Work

In this section we present a list of methods that are closely related to us namely, learning embedding on relation databases, learning global table embedding for table retrieval tasks and table augmentation.

2.1 Learning embedding on relation databases

Bordawekar and Shmueli (2017) proposed the concept of learning embeddings on relational databases by using unsupervised, distributed representation to map database entities to latent space. Their approach to capture the semantic relationship between entities in a row in the database was aimed to solve complex *cognitive* queries. Bordawekar and Shmueli (2019) further looked into novel challenges in this domain mainly addressing the degree of opacity of the disclosed database information.

2.2 Learning global table embedding for Table Retrieval tasks

Earlier approaches (Zhang and Balog, 2018) have explored the use of global table embeddings for the task of returning related tables with the representation of table and query into multiple latent semantic space and then an aggregated matching over all these spaces. Zhang and Balog (2019b) extend the above idea but instead of matching query with the table they matched tables with each other by representing both of them in latent semantic space.

2.3 Table Augmentation

Table augmentation was intended to fill missing entities in the table based on information from table corpus. Yakout et al. (2012) looked into indirectly matching web table to input using Topic Sensitive Pagerank and augmentation framework. Zhang and Chakrabarti (2013) extend the above

idea by looking into matching continuous data streams like numeric data stream and time-series data by building a semantic graph. Recently focused have been shifted to other allied fields like row and column augmentation in Zhang and Balog (2017). They focused only on those relational databases that have an entity focus. They looked into populating row with additional entries which has earlier been explored in (Wang et al. (2015), Sarma et al. (2012)) as well as populating a column with additional fields for the entity. They came up with a learning methodology for both tasks that combined information from both knowledge graph and table corpus. Table augmentation was revisited in Zhang and Balog (2019a) where they looked into probabilistic models that can be used to identify candidate value from data cell. Their approach mapped the target attribute to semantically similar column heading in table and predicate in Knowledge Base. Finally, they came up with novel Learning to Rank-based approach to rank these candidate cells.

3 Baseline Methodology

3.1 Simple Word2Vec on the Denormalized Table

Following Bordawekar and Shmueli (2017) we learn word2vec (Mikolov et al. (2013)) for entity representation with each row in the database represented as sentence and database entities as words. Then we include simple word2vec model among our baselines.

3.2 Knowledge Graph Constructed on the Denormalised Table

Here we demonstrate the procedure of creating a knowledge graph from relational databases. For each row in the database, there are multiple column entities. We create a relation between every pair of columns; for example, between column i and j , we create the relation " i_j ". Then for every row in the database, each pair of entities are joined by a relation in the knowledge graph, for example, if entity a occur in column " $i1$ " and entity b occur in column " $j1$ " in the same row in the database, then a and b get *joined* by relation " $i1_j1$ " directed from a to b and *joined* by relation " $j1_i1$ " directed from b to a .

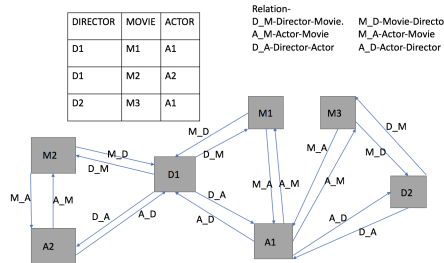


Figure 1: Creation of Knowledge Graph from a Relational Table

We then use TransH (Wang et al., 2017) to learn embeddings for entities in the knowledge graph. TransH can learn one-to-N relation which supports its suitability in the setting of a relational database.

4 Column Semantics for Denormalised Table Sampling

Databases are complicated entities and they usually consist of continuous stream of data (like integers) whereas word2vec has been traditionally applied to discrete data. In our setting, we had `movie-rank` and `director-genre` as a continuous stream of rational numbers between 0-10 and 0-1 respectively. To convert them into discrete distribution, we restrict movie rank to closest integer and director genre to the closest multiple of 0.1. However, we still would not be able to capture the inherent sequential properties in continuous streams of data. It is also important to note that in a sentence, words that are closer to each other in the sentence are more similar to each other. This approach faces a problem when we look at a row in a database since every entity in the row has *many-to-one* relation with every other entities in a row. As a result, each word's context should span the entire row. In our

experiments, we exhaustively criticize this approach when we look into the impact of changing the context window.

4.1 Genre Sampling

In our dataset, each director has multiple genres with associated director genre probability. Here it is important to note that the genre probabilities are not only continuous streams of data but also exhibit sequence and order properties. To motivate the above statement, the example where Director A has genre mystery with probability "0.6", Director B has genre mystery with probability "0.2" and Director C has genre mystery with probability "0.7". Thus, it appears that Director A should be considered to be more similar to Director C than Director B. However, it is unfair to expect from the word2vec model to learn that 0.7 and 0.6 are closer values and hence A and B are more closely related to each other. Thus, we come up with a simple fix to this problem to introduce this semantic ordering into our embeddings. We use this genre probability to sample genres for a director. We expect to make the model learn that genre associated with the director with higher probability is more closely related and hence should be sampled more times in the director's context.

We explain our methodology with an example. Suppose Director has 3 genres *A*, *B* and *C* with probabilities p_A , p_B and p_C . Then Genre *A* will be sampled with probability $\frac{p_A}{p_A+p_B+p_C}$. Genres are sampled for each director 6 times according to the above-mentioned probability. We also shuffle the entire sentence so that we can avoid changes results by a different order of joining of tables.

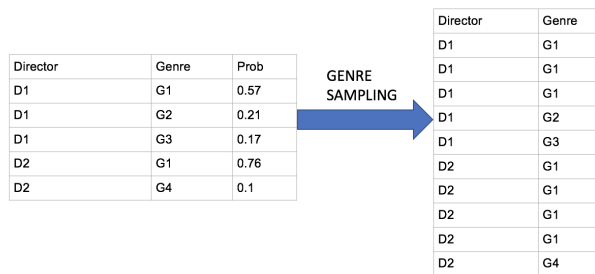


Figure 2: Analyzing Column Semantics

4.2 Movie Rank Sampling

In our data set each director has movie rank associated with each directed movie with each movie having a rank associated with it. Here for each director movies get a sample based on the movie rank. The motivation behind this variant is the fact that a director is more known for his popular films. This variant is built on top of genre sampling.

5 Incorporating inter-record Relationships

We can represent a time series for each director consisting of the movies directed by him sorted temporally. As discussed before, most of the embedding frameworks cannot effectively learn embeddings on time series data since they can not capture sequentially and thus cannot retrieve the relation that the year 1976 comes before 1999 which comes before 2002. This is important because directors who direct movies during a similar period should be similar. Moreover, the director's interest may change over time and hence we could use time series to build time-sensitive embedding for the director. But the word2vec model is incapable to learn the differences associated with these drastic changes in temporal characteristics. Thus we use time series for each director based on their movies to create "time-sensitive" director embedding based on his movie preference. We give all the movies directed by the director till time t as input time series. All the parameters of the movie at time t are provided as input at time t to LSTM. For our relational database, the parameters for each movie are actor, role, movie_genre, movie_year, and movie_rank. Since a movie has multiple values




Figure 3: LSTM for Director Embeddings

for parameter actor, role and genre, we randomly select one of these values for these parameters. Embedding of the movie is generated by concatenating actor, role, movie_genre, movie_year, and movie_rank embedding. Movie embedding at time t' is provided as input at time t' to LSTM. The LSTM output is then passed through a fully connected layer to bring it to the same dimension as movie embedding. Cosine similarity between this generated director embedding and movie embedding is used to predict the movie directed by the director at time $t+1$. We call the generated director embedding as director embedding based on movie preference.

We then optimize the director embedding by maximizing its cosine similarity with the movie directed by him at time $t+1$. To do so, we randomly sample 5 movies not directed by the director. We take the dot product of this director embedding with movie embedding of correct last movie as well as movie embedding of 5 randomly negative sampled. We then take softmax over this output obtained after dot product and optimize Negative Log-Likelihood loss to train the LSTM.

5.1 LSTM variants

Joint Loss variant It was observed that learned director embedding was close only to the movie embedding at time $t+1$ and has lost the information of all other movies directed by the director. Hence the Loss function was computed for each of the correct movies (i.e. movies used as input to LSTM as well as the movie at time $t+1$) and net loss was computed as the sum of all these losses. All other LSTM variants are built on top of Joint Loss Variants.

Actor sampling We identified that not a lot of information can be captured by only one random actor since most of the movies tend to have a large number of actors. Thus we come up with a variant where we randomly sample 3 actors and 3 roles for each movie. We then combine actor embedding for these actors based on 2 operations, one is simple aggregation by averaging whereas the other was concatenating them and passing them to linear layer to get movie's aggregated actor embedding. A similar operation was performed to generate the movie's aggregated role embedding. We compare the performance of these 2 approaches in the Results section.

Popular actor variant We first identify the popularity of an actor based on the number of movies he has worked on. We then choose the top 3 most popular actors for each movie and their corresponding roles in those movies to create movie embedding.

6 Evaluation

6.1 Data

For our experiments, we use the IMDB dataset which consists of 7 tables which when combined comprise a total of 21 columns. A table describes subtopics like director, director genre, movie, actor, etc. with columns referring to entities in this subtopic like director id, director first name, director last name for the director table.

6.2 Models

All word2vec models were trained on the entire dataset with embedding dimension 300 for 10 epochs. Just to see the effect of context window size, we experimented with a context window of size 5 for basic word2vec [Base w2v(Context 5)] whereas, for other variants, context window was taken as the size of a sentence formed from each row in the table. We also looked into the effect of the relative position of the columns. We brought the movie column next to the director column and then trained word2vec (Base w2v (Close)) just to see if there are significant changes induced by the relative position of columns. TransH models were run for 1000 iteration at a learning rate 0.001 with embedding dimension 50 for each entity and relation. For LSTM implementation, the movies directed by the director are sorted based on year. Now we pick 6 movies directed by the director- the first 5 movies are selected for the training set and the last movie for the test set. We also randomly sample 5 movies not directed by the director as negative samples for training. For each director, 80% of movies go into the training set, 10% into the validation set and the remaining 10% into the test set. Now in these 80% movies in the training set, We randomly sample 5 movies from the director's training set and 5 movies not directed by the director as negative samples for training. We repeat steps described before ($\text{len}(\text{movies in the director's training set})/5$) number of times for each director. Then first 4 movies are used as input to LSTM whereas movie 5 is used as a movie at t+1 to compute loss for LSTM. The model is trained using Adam Optimiser with a learning rate of 0.001 with a batch size of 1024. All the other variants of LSTM are also trained in similar training methodology.

6.3 Evaluation using WikiLinks

Now, the task before our hand was to find a way to evaluate our approach. We take inspiration from methodologies to obtain semantic relatedness between Wikipedia Entities. Milne Witten Inlink Score (Milne and Witten (2008)) uses the hyperlink structure of Wikipedia to measure this semantic relatedness. To evaluate our approach, we looked at the top 63 popular directors according to IMDB. We create a list of 100 most similar directors based on Milne Witten Inlink score as gold sample. We evaluate our models by generating a list of 100 most similar directors based on similarity model and comparing it with the gold sample in terms of NDCG and Precision. For NDCG computations, score 5 was given to top20 results, score 4 was given to top20-top40 results, score 3 was given to top40-top60 results, score 2 was given to top60-80 results and score 1 was given to top80-100 results and 0 to others. For precision, we assumed the top 100 most similar directors in our gold sample as relevant directors to the given director.

6.4 Evaluating Database Completion

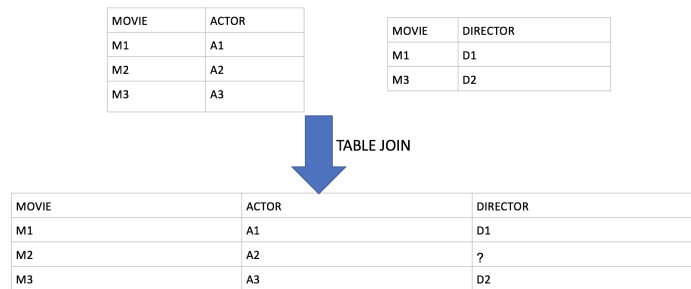


Figure 4: Envisioned Auto Completion of Data Cells in table

We randomly select 20% of movies directed by the director for the test set and put remaining movies in train set for word2vec variants. Then for every movie in the test set, We randomly sample 99 directors as negatively sampled directors. Finally, for each movie in the test set, We compute its similarity with correct as well as negatively sampled directors. We then rerank this list of candidate directors based on similarity scores and compute precision values on this reranked list.

Model	Prec@10	Prec@20	NDCG@10	NDCG@20
Base w2v (Context 5)	0.158 ^(T)	0.127	0.124 ^(T)	0.104 ^(T)
Base w2v (Close)	0.145 ^(T)	0.132	0.114 ^(T)	0.102 ^(T)
Base w2v	0.169 ^(T)	0.142 ^(T)	0.126 ^(T)	0.111 ^(T)
Genre w2v	0.196 ^(TSC)	0.161 ^(TSC)	0.144 ^(TC)	0.124 ^(Te)
Movierank w2v	0.176 ^(T)	0.149 ^(T)	0.136 ^(T)	0.120 ^(T)
TransH	0.106	0.112	0.079	0.078
LSTM loss	0.155 ^(T)	0.140	0.120 ^(T)	0.110 ^(T)

Table 1: Results on the Similarity Queries. Significance tests were conducted using a two-tailed paired Student’s t-test. Uppercase or lowercase characters in brackets indicate statistical significance with $p < 0.05$ or $p < 0.10$, respectively, over the Base w2v (B/b), Base w2v (Context 5) (S/s), Base w2v (Close) (C/c), Genre w2v (G/g), Movierank w2v (M/m), TransH (T/t) and LSTM Joint Loss Variant(L/l).

We also performed time-sensitive Database completion where instead of randomly selecting 20% of movies, we select the last 20% of movies in the test set. Then for each movie in the test set, we utilized the following methodology. If the movie was directed at time t, we use movies directed by director before time t but not more than 20 years before time t to create "time-sensitive" director embedding. We use these time-sensitive director embedding for the correct director as well as 99 negatively sampled directors and followed methodology similar to one in the paragraph before. For this, we also selected 20% movies in the validation set and the remaining 60% movies for the train set. The embeddings used in LSTM were initialized by embeddings learned from genre sampling models and were not backpropagated in our training.

7 Results

7.1 Similarity Queries

Results for solving similarity queries on databases by using wikilinks as golden standard are shown in table 1. We can see that Word2vec variants seem to be the best performing models. Genre sampling model outperforms all models in terms of NDCG and word2vec followed by MovieRank sampling models. We also see that increasing Context Window Size to the entire length of the sentence seems to improve performance. This is because in the database the columns that are adjacent to each other may not necessarily be similar to each other. Thus in a database, entities are affected by all adjoining entities in a row in the database. Also, the relative position of these entities does not affect performance as there is no statistically significant difference between them. Hence by keeping the context window large enough we can overcome the effect of the relative position of columns. Hence we just shuffle row before performing word2vec in Genre Sampling and Movie Rank Sampling models. We can conclude from this table that sampling based on integer values may not always produce the most idealistic results. In this case, movies that might not be popular can still give crucial insight that could be used as a distinctive feature of a director like what genre of movies are preferred by the director and what actors do he usually likes to work with.

It was also observed that about half of the directors in our database were not Wikipedia entities. This might adversely taint our results since there is a possibility that there exists a director that is very similar to our candidate director but it does not happen to be a Wikipedia entity. Thus we recompute the result for similarity queries but this time reranking only those directors that are in the Wikipedia database. The results can be seen in table 2. The trends are similar to as seen before but the Precision and NDCG values are much better however we can see that Genre Sampling and Movie Rank Sampling variants of word2vec are statistically significantly better than word2vec and TransH baseline models.

We also tried to evaluate if there is any difference in result over the performance of these semantic queries if we re-rank only popular directors i.e. those that have directed at least 5 movies. We can see the results in table 3. We see that here MovieRank Sampling performs better than genre sampling. We attribute this to the reasoning that for popular directors, the similarity between popular movies serve as an important factor for semantic similarity. Other trends are similar to as seen above.


Model	Prec@10	Prec@20	NDCG@10	NDCG@20
Base w2v (Context 5)	0.182	0.155	0.142	0.123
Base w2v (Close)	0.153	0.139	0.122	0.108
Base w2v	0.188	0.161	0.142	0.127
Genre w2v	0.248 (TSBC)	0.208 (TSBC)	0.186 (TSBC)	0.162 (TSBC)
Movierank w2v	0.235 (TSBC)	0.2 (TSBC)	0.187 (TSBC)	0.164 (TSBC)
TransH	0.161	0.15	0.115	0.106
LSTM loss	0.192	0.165	0.146	0.132

Table 2: Results on the Similarity Queries for reranking only those entities in Wikipedia. Significant test performed the same as before.

Model	Prec@10	Prec@20	NDCG@10	NDCG@20
Base w2v	0.202 (T)	0.175	0.150 (T)	0.134 (T)
Genre w2v	0.229 (T)	0.210 (TIB)	0.174 (T)	0.165 (TB)
Movierank w2v	0.242 (TB)	0.218 (TLB)	0.181 (TB)	0.164 (TB)
TransH	0.145	0.154	0.101	0.103
LSTM loss	0.194	0.173	0.153 (t)	0.137

Table 3: Results on the Similarity Queries for reranking only those directors that have directed at least 5 movies. A significant test performed the same as before.

7.2 Database Completion



(a) Database Completion Results Random Removal (b) Database Completion Results Time Specific

Figure 5: DataBase Completion Results

Now we look into results of Database Completion for word2vec variants in 5. We see that again Genre Sampling and MovieRank Sampling outperform the basic Word2Vec model. Moreover, we see a very positive performance of nearly 0.8 for precision@20 providing evidence that our approach can be successfully utilized in the auto-completion of data cells. In the second graph, we do time-specific database completion. We see that LSTM variants perform better than simple Joint Loss LSTM. LSTM Sampling and LSTM popular sampling variant model seems to be best performing models with both achieving around 0.6 Precision@20.

8 Conclusions

Learning entity embeddings in a relational database presents interesting challenges over traditional models over text data or knowledge graphs due to various complex data semantics that co-exist. We believe that our results, though preliminary, demonstrate not only the need for embedding entities in relational databases, but also to guide these embeddings carefully by using data semantics. Going further, we plan to investigate how to formally share these data semantics declaratively with representation learning frameworks over relational databases.

References

- Imdb dataset. URL <https://relational.t.cvut.cz/dataset/IMDb>.
- R. Bordawekar and O. Shmueli. Using word embedding to enable semantic queries in relational databases. In *Proceedings of the Data Management for End to End Learning Work Workshop, SIGMOD'17*, 2017.
- R. Bordawekar and O. Shmueli. Exploiting latent information in relational databases via word embedding and application to degrees of disclosure. In *Proceedings of 9th Biennial Conference on Innovative Data Systems Research(CIDR'19)*, CIDR'19, 2019.
- J. Chung, K. Kastner, L. Dinh, K. Goel, A. Courville, and Y. Bengio. A recurrent latent variable model for sequential data. In *NIPS*, 2016.
- S. Hochreiter and J. Schmidhuber. Long short-term memory. In *Neural Comput.*, 1997.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- Y. Lin, Z. Liu, M. Sun, Y. Liu, and X. Zhu. Learning entity and relation embeddings for knowledge graph completion. In *AAAI*, 2015.
- T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. In *Proceedings of Twenty-seventh Conference on Neural Information Processing Systems, NIPS'13*, 2013.
- D. Milne and I.H. Witten. An effective, low-cost measure of semantic relatedness obtained from wikipedia links. In *Proceedings of the Nineteenth AAAI Conference on Artificial Intelligence, AAAI'08*, 2008.
- A.D. Sarma, L. Fang, N. Gupta, A. Halevy, H. Lee, F. Wu, R. Xin, and C. Yu. Finding related tables. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data, SIGMOD '12*, 2012.
- C. Wang, K. Chakrabarti, Y. He, K. Ganjam, Z. Chen, and P.A. Bernstein. Concept expansion using web tables. In *Proceedings of the 24th International Conference on World Wide Web, WWW'15*, 2015.
- Z. Wang, J. Zhang, J. Feng, and Z. Chen. Knowledge graph embedding by translating on hyperplanes. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, AAAI'17*, 2017.
- Ian H Witten and David N Milne. An effective, low-cost measure of semantic relatedness obtained from wikipedia links. 2008.
- M. Yakout, K. Ganjam, K. Chakrabarti, and S. Chaudhuri. Infogather: Entity augmentation and attribute discovery by holistic matching with web tables. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data, SIGMOD'12*, 2012.
- M. Zhang and K. Chakrabarti. Infogather+: Semantic matching and annotation of numeric and time-varying attributes in web tables. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data, SIGMOD'13*, 2013.
- S. Zhang and K. Balog. Entitables: Smart assistance for entity-focused tables. In *Proceedings of the 40nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'17*, 2017.
- S. Zhang and K. Balog. Ad hoc table retrieval using semantic similarity. In *Proceedings of the The 2018 Web Conference, WWW'18*, 2018.
- S. Zhang and K. Balog. Auto-completion for data cells in relational tables. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM'19*, 2019a.
- S. Zhang and K. Balog. Recommending related tables. In *CoRR*, 2019b.