Efficient Taxonomic Similarity Joins with Adaptive Overlap Constraint

Pengfei Xu and Jiaheng Lu
Department of Computer Science, University of Helsinki
first.last@helsinki.fi

ABSTRACT

A similarity join aims to find all similar pairs between two collections of records. Established approaches usually deal with synthetic differences like typos and abbreviations, but neglect the semantic relations between words. Such relations, however, are helpful for obtaining high-quality joining results. In this paper, we leverage the taxonomy knowledge (i.e., a set of IS-A hierarchical relations) to define a similarity measure which finds semantic-similar records from two datasets. Based on this measure, we develop a similarity join algorithm with prefix filtering framework to prune away irrelevant pairs effectively. Our technical contribution here is an algorithm that judiciously selects critical parameters in a prefix filter to maximise its filtering power, supported by an estimation technique and Monte Carlo simulation process. Empirical experiments show that our proposed methods exhibit high efficiency and scalability, outperforming the state-of-art by a large margin.

CCS CONCEPTS

• **Information systems** → **Join algorithms**; **Data cleaning**; *Semi-structured data*; *Inconsistent data*; *Entity resolution*;

KEYWORDS

Similarity join, taxonomic similarity, prefix filtering, estimation

1 INTRODUCTION

Given two sets of records, a similarity join aims to find all records whose similarities are higher than a given threshold. Such operation is widely-seen in tasks such as data cleaning [1, 7], information retrieval [2, 13], and data mining [6]. To perform joining efficiently, a plethora of established algorithms utilise similarity measures, e.g., *Levenshtein similarity* [11] and *Jaccard coefficient* [9]. Such measures capture syntactic-similar records, which is not enough because of the existence of synonyms and related concepts, which often differ from spellings.
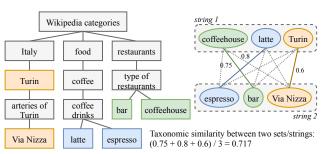
Figure 1: Example of a simplified hierarchical taxonomy and similarity calculation based on Equations 1 and 2.

Taxonomy is an abundant source of lexica, maintaining IS-A relations between terms. It has been proved useful [10] for enhancing the quality of similarity joins. Figure 1 depicts an example. Given a Wikipedia taxonomy and two strings, a join algorithm with Levenshtein distance will fail to capture their similarity due to distinct spellings. In contrast, a *taxonomic similarity measure* maps each string to multiple taxonomy nodes, and calculate the similarity between every two nodes from the depth of their lowest common ancestor (LCA). For example, "Turin" and "Via Nizza", with three common ancestors (including "Turin" itself) and maximal depth five, have 0.6 (=3/5) similarity. Hence, the similarity between two strings can be calculated as 0.717, by averaging the maximal sum of three distinct node-wise similarities.

Joins with taxonomic similarity measure can be useful in many real-life scenarios. For example, location providers are interested in integrating knowledge taxonomy to remove duplicates or link related records from crawled Points of Interests (POIs) [10]; *personalised medicine* provides specific treatment to a small group of patients clustered by a disease taxonomy [4]. Also, taxonomic joins can be helpful for enhancing the quality of similarity matrices used in various recommender systems [14].

In this article, we tackle the taxonomic joining problem by adopting the *filtering-and-verification* framework, which works by first (i) removing record pairs which are impossible to be similar, then (ii) verifying the real similarity of survived pairs. Since verification is expensive, we introduce a novel filtering technique which efficiently removes unfeasible pairs whose number of similar nodes do not reach a given *overlap constraint*. We observed that such constraint vitally affects filtering quality and thus joining time. In this paper, we propose an estimator to predict the running time of a given constraint by performing test-drives on small samples and then scaling the result accordingly. To suggest the best constraint for minimising the joining time, as our technical contribution, we propose a Monte Carlo simulation process which gives accurate suggestions without a predetermined sample size. Experiments show

that, remarkably, our suggestion achieves higher than 90% accuracy, by using only a few (e.g., 100) samples from nearly 3 million records, and occupies approximately 1% of total joining time. The superior performance of our solution over the state-of-art approach [10] further motives its application in practice.

2 PRELIMINARIES

Similarity measures. Let $S : \{s_1, \cdots, s_i\}$ and $T : \{t_1, \cdots, t_j\}$ be two sets of nodes from a hierarchical taxonomy. Let $s \in S$ and $t \in T$ be two nodes, and let $|s|$ ($|t|$) denotes the depth of node s (t). Then, their similarity can be measured based on the depth of the lowest common ancestor (LCA):

$$TS(s, t) = \frac{|LCA(s, t)|}{\max(|s|, |t|)} \tag{1}$$

With Equation 1, the similarity between two sets S and T can then be obtained by averaging the maximum sum of all TS's of *distinct* node pairs, where $|S|$ ($|T|$) is the number of nodes in set S (T), I_{pq} is an indicator variable (i) controlling whether to select the edge (s_p, t_q), and (ii) ensures any of s_p or t_q is used at most once:

$$GTS(S, T) = \frac{W(S, T)}{\max(|S|, |T|)} = \frac{\max \sum_p \sum_q I_{pq} TS(s_p, t_q)}{\max(|S|, |T|)} \tag{2}$$

where $p \in [1, i], q \in [1, j], I_{pq} = 0$ or $1, \sum_p I_{pq} \leqslant 1,$ and $\sum_q I_{pq} \leqslant 1$

Solving for the value of W in Equation 2 requires to find the maximum weight matching in a bipartite graph, which can be categorised as an *assignment problem. Hungarian algorithm* [8] is so far the best solution which runs in a polynomial $O(n^3)$ time.

Example 2.1. Take two strings in Figure 1 as an example. Since the three most-similar node pairs are ("coffeehouse", "bar"), ("latte", "espresso"), and ("Turin", "Via Nizza"), the GTS similarity between two strings becomes 0.717 (= (0.75 + 0.8 + 0.6)/3). Note that the distinctness forbid any node from being selected more than once, e.g., selecting both ("latte", "espresso") and ("latte", "Turin") are not allowed.

Problem definition. We define our research problem as follows:

PROBLEM 1. *Let $\mathcal{S}$ and $\mathcal{T}$ be two collections of sets $\mathcal{S} : \{S_1, \cdots, S_m\},$ $\mathcal{T} : \{T_1, \cdots, T_n\},$ where each set contains multiple nodes, i.e., $S \in \mathcal{S} : \{s_1, \cdots, s_i\}, T \in \mathcal{T} : \{t_1, \cdots, t_j\}.$ Given a GTS similarity measure and a similarity threshold $\theta,$ find all pairs of sets in forms of $(S, T) \in \mathcal{S} \times \mathcal{T}$ such that each $GTS(S, T) \geqslant \theta.$*

It is not trivial to solve Problem 1 efficiently. Recall Example 2.1. To apply the Hungarian algorithm, we first need to fill a 3×3 matrix by 9 TS calculations, not to mention a longer string which may have hundreds of words resulting in a massive amount of calculations. Large datasets exacerbate the situation as processing every string pair leads to interminable running time. Hence, it is crucial to have an efficient solution which avoids running Hungarian algorithm whenever possible to speed up the joining process.

3 ADAPTIVE OVERLAP JOINING

We now present our novel joining algorithm which include three stages, namely (i) *inverted lists construction* (Line 2 in Alg. 1), where each set is being indexed for faster overlap-finding; (ii) *filtering* (Lines 3 - 14), where we try to purge unfeasible set pairs which

have not enough ($\geqslant \tau$) similar nodes; and (iii) *verification* (Lines 15 - 16) where we perform actual *GTS* calculation on survived pairs. Since Stage 2 is a key step to speed up the whole processing, we will focus on developing an effective optimisation strategy.

Based on the definition of *GTS*, our filtering technique states that two similar sets must have at least τ pairs of similar nodes, where τ is the overlap constraint:

LEMMA 3.1 (AP-FILTER). *Given two sets S and T, and without loss of generality by assuming $|S| < |T|$. If $GTS(S, T) \geqslant \theta$, then there are at least τ distinct similar pairs of nodes $(s_i \in S, t_j \in T)$ such that each of them satisfies $TS(s_i, t_j) \geqslant \varphi$, where $\varphi = \frac{\theta|T| - \tau + 1}{|S| - \tau + 1}$.*
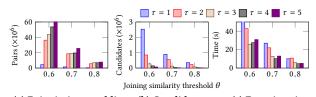
According to Lemma 3.1, for each pair of sets, we need to find the number of distinct node pairs whose $TS \geqslant \varphi$. Since TS depends on $|LCA|$, the problem can be converted to overlap finding problem, which is to find the common ancestors between pairs of nodes within two sets. An efficient way to find such ancestors is to index them (as keys, corresponding nodes as values) by *inverted lists*, and joining these lists afterwards to obtain overlapped keys and thus sets which contains nodes having common ancestors. Furthermore, Lemma 3.1 can be relaxed to the follows, so that we can build an inverted list independently for each of $\mathcal{S}$ and $\mathcal{T}$:

COROLLARY 3.2. *Two similar sets, S and T, must have at least τ distinct (s_i, t_j)'s such that each $TS(s_i, t_j) \geqslant \varphi$, where $\varphi = \frac{\theta|S| - \tau + 1}{|S| - \tau + 1}$.*

Corollary 3.2 states that a node $s \in S$ only needs to put its ancestors deeper than $\varphi|s|$ into the inverted list since others are not able to achieve φ similarity. This insight reduces the index size and thus accelerate filtering.

The pseudo code of our joining technique is presented in Algorithm 1. Specifically, Line 7 employs a *length filtering* technique to remove pairs having a disparate number of nodes, and Line 10 ensures the distinctness of each node so that a node s is counted one time even if it is similar to multiple t's from another set T.

Algorithm 1: AP-JOIN: Set joining with AP-FILTER

Input: two collections of sets $\mathcal{S}$ and $\mathcal{T}$, a similarity threshold θ and an positive integer τ

Output: R: $\{(S, T) \in \mathcal{S} \times \mathcal{T} | GTS(S, T) \geqslant \theta\}$

1   $\mathcal{P} \leftarrow \varnothing, C \leftarrow \varnothing, \mathcal{R} \leftarrow \varnothing$

2   $\mathcal{L}_\mathcal{S}, \mathcal{L}_\mathcal{T} \leftarrow$ inverted lists built from $\mathcal{S}$ and $\mathcal{T}$, according to Cor. 3.2

3   $\mathcal{G} \leftarrow$ overlapped keys (common ancestors) between $\mathcal{L}_\mathcal{S}$ and $\mathcal{L}_\mathcal{T}$

4   **foreach** $g \in \mathcal{G}$ **do**

5     $\ell_\mathcal{S}, \ell_\mathcal{T} \leftarrow$ lists indexed by g in $\mathcal{L}_\mathcal{S}$ and $\mathcal{L}_\mathcal{T}$

6   **foreach** $(S, T) \in (\ell_\mathcal{S} \times \ell_\mathcal{T})$ **do**

7   **if** $\min(|S|, |T|) \leqslant \theta \cdot \max(|S|, |T|)$ **then**

8   **continue**   // length filtering

9       $s_g, t_g \leftarrow$ nodes producing ancestor g in S and T

10   **if** s_g and t_g have not been found similar with other nodes from T or S **then** $\mathcal{P} \leftarrow \mathcal{P} \cup \{(S, T)\}$

11   **foreach** $(S, T) \in \mathcal{P}$ **do**   // find pairs with τ overlaps

12   **if** (S, T) appears at least τ times in $\mathcal{P}$ **then**

13       $C \leftarrow C \cup \{(S, T)\}$

14   **foreach** $(S, T) \in C$ **do**   // verification

15   **if** $GTS(S, T) \geqslant \theta$ **then** $\mathcal{R} \leftarrow \mathcal{R} \cup \{(S, T)\}$

16   **return** $\mathcal{R}$

(a) Pairs in inverted lists   (b) Candidates   (c) Running time

Figure 2: Overlap constraint τ's affecting joining performance ($10K \times 10K$ subsets of OHSUMED).

Effect of overlapping constraint τ. The value of parameter τ influences the joining time. Intuitively, as τ increases, the sizes of inverted lists will grow, indicating that the *filtering* will be slower. In the meantime, fewer pairs satisfies the increased overlap constraint, leading to the faster *verification* phase. Perceiving the opposite trends, we conduct an empirical experiment as presented in Figure 2, in which we confirm the existence of optimum that minimises the overall running time. Now a natural question arises: *how to find such optimal τ correctly and efficiently?* We tackle the question and present our answer in the next section.

4 PARAMETER RECOMMENDATION

This section aims to give an accurate recommendation for τ which minimises the total joining time. It is backed by a cost model and a sampling-based estimator.

Cost model. The joining time cost can be modelled as follows:

$$C_\tau = C_{F_\tau} + C_{V_\tau} = t_F \cdot F_\tau + t_V \cdot V_\tau \tag{3}$$

where the total cost C_τ is the sum of filtering (C_{F_τ}), and verification cost (C_{V_τ}). Each of them is obtained by multiplying corresponding number of processed pairs (F_τ or V_τ) by the average time to process one pair in each stage (t_F or t_V).

Bernoulli estimator. It is certainly unfeasible to run the joining algorithm on full datasets to get its cost. Instead, we can *estimate* the number of pairs in each stage ($\hat{F}_\tau$ and $\hat{V}_\tau$) and hence the total cost $\hat{C}_\tau$ by using the *independent Bernoulli sampling*, where each set in input dataset $\mathcal{S}$ ($\mathcal{T}$) has probability p_s (p_t) for being in the sample. Therefore, a set pair (S, T) being processed during a *real* filtering (or verification) stage has probability $p_s p_t$ for being counted into F_τ (or V_τ), i.e., when once both S and T exist in the sample. Hence, we get an unbiased estimator of F_τ:

$$E[F'_\tau] = F_\tau \cdot p_s p_t \Rightarrow \hat{F}_\tau = \frac{F'_\tau}{p_s p_t}. \text{ Similarly, } \hat{V}_\tau = \frac{V'_\tau}{p_s p_t} \tag{4}$$

Plugging $\hat{F}_\tau$ and $\hat{V}_\tau$ into Equation 3 to obtain estimated cost $\hat{C}_\tau$.

Iterative suggestion refinement. Equation 4 is a static estimation strategy where p_s and p_t are determined beforehand, usually by trial-and-error. To erase the requirement of this foreknowledge, we propose an iterative method based on Monte Carlo simulation, which refines the suggestion from multiple iterations until the smallest $\hat{C}_\tau$ is identified with high confidence.

Multiple iterations give a series of estimations. Since all of them are i.i.d., the Central Limit Theorem (CLT) holds such that their mean $\hat{\mu}_{F_\tau}$ converges to a normal distribution when iteration goes on (the same for $\hat{\mu}_{V_\tau}$; $Var[\cdot]$ denotes the population variance):

$$\hat{\mu}_{F_\tau} \sim \mathcal{N}\big(E[\hat{F}_\tau], Var[\hat{F}_\tau]/n\big) \text{ when } n \to \infty \tag{5}$$
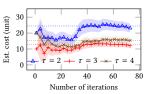

Figure 3: Illustration of the convergent of the means (solid lines) and CI's (shaded areas) to the real values (dotted lines). The cost when $\tau = 1$ is too large thus have been eliminated.

CLT also allude that the mean and variance of underlying distribution, $E[\hat{F}_\tau]$ and $Var[\hat{F}_\tau]/n$, can be estimated by the sample mean and variance, $\hat{\mu}_{F_\tau}$ and $\hat{\sigma}^2_{F_\tau}$. Both estimators are unbiased. Thus, we can calculate $\hat{\mu}_{F_\tau}$ and $\hat{\sigma}^2_{F_\tau}$ by a recursive formula (e.g., [3]), and use them to estimate $\hat{\mu}_{F_\tau}$ (the same for $\hat{\mu}_{V_\tau}$). Since both $\hat{F}_\tau$ and $\hat{V}_\tau$ converge to normal distributions, the estimated total cost $\hat{C}_\tau$ also converges, and can be modelled by a *Student's t distribution*. The confidence interval (CI) of C_τ can be constructed consequently.

Stopping criterion. Given a universe of τ's, we can safely terminate the refinement procedure once the overlapped CI's between the best (which gives the least $\hat{\mu}_{C_\tau}$) and other τ's are small enough. In other words, the refinement stops when *the additional cost due to an inaccurate estimation (i.e., the sum of overlapped CI's) is less than that for one more iteration*:

LEMMA 4.1 (STOPPING CRITERION). *Let $\mathbb{U}_\tau$ be the universe of τ's, and let $\tau_1 \in \mathbb{U}_\tau$ denote the τ leading to the minimal estimated cost, i.e., $\tau_1 = \arg\min_\tau (\hat{\mu}_{C_\tau})$. The refinement process terminates when $\sum_{\tau_2} (U_{C_{\tau_1}} - L_{C_{\tau_2}}) < t_F \cdot \sum_\tau F'_\tau$ holds for all $\tau_2 \in \mathbb{U}_\tau, \tau_2 \neq \tau_1$.*

We present our iterative refinement procedure in Algorithm 2. Given multiple independent samples form $\mathcal{S}$ and $\mathcal{T}$, it runs the filtering stage of AP-JOIN for every τ, obtaining F'_τ and V'_τ, then estimates the mean and variance of $\hat{C}_\tau$. The procedure terminates when the best τ is found with a predefined confidence level. The refinement runs at least n^* iterations to discard the effect of instability in the early stage, known as the *burn-in* period.

5 EXPERIMENTAL ANALYSIS

We implemented all algorithms in Java 8, and run the code on a quad-core Xeon 2.53GHz node with 32GB RAM. We use two

---

Algorithm 2: Cost-bounded suggestion refinement

Input: k samples from each dataset: $\{\mathcal{S}'_1, \cdots, \mathcal{S}'_k\}$ and $\{\mathcal{T}'_1, \cdots, \mathcal{T}'_k\}$, a positive integer n^*, a Student's t quantile t_* corresponding to a specific confidence level, and a universe of τ's $\mathbb{U}_\tau$

Output: τ which has the minimal estimated cost

1   $n \leftarrow 1$ // select first samples
2   **repeat**
3   **foreach** $\tau \in \mathbb{U}_\tau$ **do**
4   Run the filtering stage of AP-JOIN on samples $\mathcal{S}'_n$ and $\mathcal{T}'_n$
5   Compute the confidence interval of total cost $CI(C_\tau)$
6      $n \leftarrow n + 1$ // select next samples
7   **until** $n \geqslant n^*$ *and the condition in Lemma 4.1 is fulfilled*
8   **return** τ_1 *in Lemma 4.1*

Table 1: Characteristics of taxonomies and datasets.

Taxonomy trees (Height in min/avg/max)				String datasets (# nodes in min/avg/max)		
Name	# of nodes	Height	Avg. fanout	Name	# of records	# of nodes
MeSH	57,840	1 / 5.1 / 12	157	OHSUMED	293,294	5 / 8.4 / 26
Wiki categories	1,212,943	1 / 6.2 / 26	32,300	Wiki articles	3,512,954	5 / 8.2 / 277

Table 2: Performance vs the state-of-art w.r.t. θ.

Dataset	Algorithm	# Pairs (10^8)			# Candidates (10^6)			Running time (min)		
		0.6	0.7	0.8	0.6	0.7	0.8	0.6	0.7	0.8
Wiki articles (50K×50K)	AP-Join	**0.42**	**0.08**	**0.01**	**11.52**	**3.04**	**0.27**	**10.03**	**2.64**	**0.55**
	K-Join	0.87	0.25	0.07	28.42	8.35	1.98	22.28	6.65	1.74
OHSUMED (50K×50K)	AP-Join	**1.08**	4.97	1.72	**63.43**	**0.64**	**0.26**	**41.81**	**4.44**	**1.67**
	K-Join	-	2.13	0.86	-	115.58	38.42	-	80.01	25.33

Table 3: Suggestion accuracy and speed (100 samples).

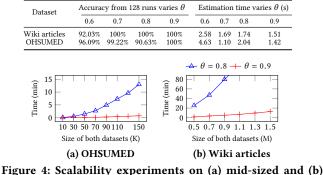| Dataset	Accuracy from 128 runs varies θ				Estimation time varies θ (s)			
	0.6	0.7	0.8	0.9	0.6	0.7	0.8	0.9
Wiki articles	92.03%	100%	100%	100%	2.58	1.69	1.74	1.51
OHSUMED	96.09%	99.22%	90.63%	100%	4.63	1.10	2.04	1.42

(a) OHSUMED   **(b) Wiki articles**

Figure 4: Scalability experiments on (a) mid-sized and (b) large-sized datasets. Two joining datasets have the same size.

taxonomies, *Wiki categories*[1] and *MeSH terms*[2], and two string datasets, *Wiki articles* and *OHSUMED articles*. Each string is mapped to a set of taxonomic nodes according to its categories (Wiki articles) or keywords (OHSUMED). Table 1 describes these datasets.

Performance vs the state-of-art. We obtained the source code from the authors of the state-of-art approach K-Join [10], rewrote their C++ code using Java, and extended the algorithm to perform R-S join following their instructions. After that, we ran both AP-Join and K-Join to compare their performance and present Table 2. It shows that our approach outperforms the state-of-art on both datasets, especially OHSUMED, with a large margin. Specifically, given the threshold 0.6, K-Join uses up 32GB of RAM for storing all pairs and eventually crashed. AP-Join, on the contrary, has fewer candidates and successfully finishes the joining.

Scalability. We randomly sample our datasets into different subsets so that the largest one contains roughly half of total records. Then, we run our algorithm and present the result in Figure 4. The result shows that the joining time increases linearly with data size. Besides, a larger θ leads to a faster joining, because a high similarity threshold leads to a high φ, which reduces sizes of inverted lists, the number of candidates, and ultimately the total joining time.

Parameter suggestion. The final experiment is the iterative procedure which suggests the key parameter τ. We set $n^* = 10, t_* = 1.036$ (70% confidence level) following common practices, and the sample size be 100, a tiny fraction comparing to whole datasets. We repeat our suggestion algorithm for 128 times and record the number of correct suggestions according to our empirical knowledge. The results (in Table 3) show that our algorithm suggests the correct τ

[1]http://wiki.dbpedia.org/services-resources/documentation/datasets
[2]https://www.nlm.nih.gov/mesh and http://trec.nist.gov/data/t9_filtering.html

values in few seconds, and achieves higher than 90% accuracy with only 100 samples for each iteration.

6 RELATED WORK

Most of recent works on set-similarity joins (e.g., [6, 10]) follow the filter-and-verification framework. Verification is expensive; hence, the key technical challenge is to design a filtering mechanism to prune away irrelevant record pairs as much as possible. Several techniques have been proposed, such as length [5], position [12], and prefix filtering [2]. While the last one is widely used in many works of literature, it suffers the problem of numerous candidates since it finds all record pairs having at least one overlapped token. This problem is tackled by Wang *et al.* [11] with a time-complexity-based method. For our research problem, the most recent article [10] extends the prefix filtering by considering the weight of each token, but still having the same problem due to the one-overlap policy (see $\tau = 1$ in Figure 2). In contrast, our work finds the best overlap constraint for each dataset to achieve a much shorter joining time.

7 CONCLUSION AND FUTURE WORK

This paper studies a problem of integrating taxonomies for efficient set-similarity joins. We first extend the prefix filtering technique to solve the join problem efficiently, then, as our technical contribution, we propose a novel estimation framework to judiciously select the parameter which minimises the total running time. Experiments based on real datasets exhibit the superiority of proposed algorithms. As future work, we would like to apply our estimation method for accelerating the computation of knowledge-based similarity matrices used in various machine learning tasks (e.g., [14]).

REFERENCES

[1] Arvind Arasu, Surajit Chaudhuri, and Raghav Kaushik. 2008. Transformation-based Framework for Record Matching. In *ICDE*. 40–49.
[2] Roberto J Bayardo, Yiming Ma, and Ramakrishnan Srikant. 2007. Scaling up all pairs similarity search. In *WWW*. ACM, 131–140.
[3] Tony Finch. 2009. Incremental calculation of weighted mean and variance. *University of Cambridge* 4 (2009), 11–15.
[4] Gavin Giovannoni. 2017. Personalized medicine in multiple sclerosis. *NDM* 7, 6s (2017), 13–17.
[5] Chen Li, Jiaheng Lu, and Yiming Lu. 2008. Efficient Merging and Filtering Algorithms for Approximate String Searches. In *ICDE*. 257–266.
[6] Jiaheng Lu, Chunbin Lin, Wei Wang, Chen Li, and Haiyong Wang. 2013. String similarity measures and joins with synonyms. In *SIGMOD*. ACM, 373–384.
[7] Jiaheng Lu, Chunbin Lin, Wei Wang, Chen Li, and Xiaokui Xiao. 2015. Boosting the Quality of Approximate String Matching by Synonyms. *ACM Trans. Database Syst.* 40, 3 (2015), 15:1–15:42.
[8] J Munkres. 1957. Algorithms for the Assignment and Transportation Problems. *JSIAM* 5, 1 (1957), 32–38.
[9] Sunita Sarawagi and Alok Kirpal. 2004. Efficient set joins on similarity predicates. In *SIGMOD*. 743–754.
[10] Zeyuan Shang, Yaxiao Liu, Guoliang Li, and Jianhua Feng. 2016. K-Join: Knowledge-Aware Similarity Join. *TKDE* 28, 12 (2016), 3293–3308.
[11] Jiannan Wang, Guoliang Li, and Jianhua Feng. 2012. Can we beat the prefix filtering?: an adaptive framework for similarity join and search. In *SIGMOD*. 85–96.
[12] Chuan Xiao, Wei Wang, Xuemin Lin, and Jeffrey Xu Yu. 2008. Efficient similarity joins for near duplicate detection. In *WWW*. ACM, 131–140.
[13] Pengfei Xu and Jiaheng Lu. 2017. Top-k String Auto-Completion with Synonyms. In *DASFAA*. 202–218.
[14] Yongfeng Zhang, Guokun Lai, Min Zhang, Yi Zhang, Yiqun Liu, and Shaoping Ma. 2014. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In *SIGIR*. ACM, 83–92.