

On the Efficient Execution of Bounded Jaro-Winkler Distances

Editor(s): Michelle Cheatham, Wright State University, USA; Isabel F. Cruz, University of Illinois at Chicago, USA; Jérôme Euzenat, INRIA & University of Grenoble, France; Catia Pesquita, University of Lisbon, Portugal

Solicited review(s): Two anonymous reviewers

Kevin Dreßler^a, Axel-Cyrille Ngonga Ngomo^a

^a*AKSW Research Group, University of Leipzig, Germany*

Abstract. Over the last years, time-efficient approaches for the discovery of links between knowledge bases have been regarded as a key requirement towards implementing the idea of a Data Web. Thus, efficient and effective measures for comparing the labels of resources are central to facilitate the discovery of links between datasets on the Web of Data as well as their integration and fusion. We present a novel time-efficient implementation of filters that allow for the efficient execution of bounded Jaro-Winkler measures. We evaluate our approach on several datasets derived from DBpedia 3.9 and LinkedGeoData and containing up to 10^6 strings and show that it scales linearly with the size of the data for large thresholds. Moreover, we also show that our approach can be easily implemented in parallel. We also evaluate our approach against SILK and show that we outperform it even on small datasets.

Keywords: Link Discovery, String similarity measure

1. Introduction

The Linked Open Data Cloud (LOD Cloud) has developed to a compendium of more than 2000 datasets over the last few years.¹ For example, data sets pertaining to more than 14 million persons have already been made available on the Linked Data Web.² While this number is impressive on its own, it is well known that the population of the planet has surpassed 7 billion people. Hence, the Web of Data contains information on less than 1% of the overall population of the planet (counting both the living and the dead). The output of open-government movements³, scientific conferences⁴, health data⁵ and similar endeavours

yet promises to make massive amounts of data pertaining to persons available in the near future. Dealing with this upcoming increase of the number of resources on the Web of data requires providing means to integrate these datasets with the aim to facilitate statistical analysis, data mining, personalization, etc. However, while the number of datasets on the Linked Data Web grows drastically, the number of links between datasets still stagnates.⁶ Addressing this lack of links requires solving two main problems: the quadratic time complexity of link discovery (efficiency) and the automatic support of the detection of link specifications (effectiveness). In this paper, we address the efficiency of the execution of bounded Jaro-Winkler measures,⁷ which are known to be effective when comparing person names [11]. To this end, we derive equations that allow discarding a large number of computations while ex-

¹See <http://stats.lod2.eu> for an overview of the current state of the Cloud. Last access: July 11th, 2014.

²Data collected from <http://stats.lod2.eu>. Last access: July 11th, 2014.

³See for example <http://data.gov.uk/>.

⁴See for example <http://data.semanticweb.org/>

⁵<http://aksw.org/Projects/GHO>

⁶<http://linklion.org>

⁷We use bounded measures in the same sense as [18], i.e., to mean that we are only interested in pairs of strings whose similarity is greater than or equal to a given lower bound.

ecuting bounded Jaro-Winkler comparisons with high thresholds.

The contributions of this paper are as follows:

1. We derive length- and range-based filters that allow reducing the number of strings t that are compared with a string s .
2. We present a character-based filter that allows detecting whether two strings s and t share enough resemblance to be similar according to the Jaro-Winkler measure.
3. We evaluate our approach w.r.t. to its runtime and its scalability with several threshold settings and dataset sizes.

The rest of this paper is structured as follows: In Section 2, we present the problem we tackled as well as the formal notation necessary to understand this work. In the subsequent Section 3, we present the three approaches we developed to reduce the runtime of bounded Jaro-Winkler computations. We then evaluate our approach in Section 4. Related work is presented in Section 5, where we focus on approaches that aim to improve the time-efficiency of link discovery. We conclude in Section 6. The approach presented herein is now an integral part of LIMES.⁸

This paper is an extended version of [6]. We improved the scalability of the approach presented in the previous version of the paper by using tries. In addition, we provide a parallel version of the approach for additional scalability. Finally, an extended evaluation of the algorithm is presented.

2. Preliminaries

In the following, we present some of the symbols and terms used within this work.

2.1. Link Discovery

In this work, we use *link discovery* as a hypernym for deduplication, record linkage, entity resolution and similar terms used across literature. The formal specification of link discovery adopted herein is tantamount to the definition proposed in [17]: Given a set S of source resources, a set T of target resources and a relation R , our goal is to find the set $M \subseteq S \times T$ of pairs (s, t) such that $R(s, t)$. If R is `owl:sameAs`, then we are faced with a *deduplication task*. Given that the ex-

PLICIT computation of M is usually a very complex endeavour, M is most commonly approximated by a set $M' = \{(s, t, \delta(s, t)) \in S \times T \times \mathbb{R}^+ : \delta(s, t) \geq \theta\}$, where δ is a (potentially complex) similarity function and $\theta \in [0, 1]$ is a similarity threshold. Given that this problem is in $O(n^2)$, using naïve algorithms to compare large S and T is most commonly impracticable. Thus, time-efficient approaches for the computation of bounded measures have been developed over the last years for measures such as the Levenshtein distance, Minkowski distances, trigrams and many more [16].

In this paper, we thus study the following problem: Given a threshold $\theta \in [0, 1]$ and two sets of strings S and T , compute the set $M' = \{(s, t, \delta(s, t)) \in S \times T \times \mathbb{R}^+ : \delta(s, t) \geq \theta\}$. Two categories of approaches can be considered to improve the runtime of measures: Lossy approaches return a subset M'' of M' which can be calculated efficiently but for which there are no guarantees that $M'' = M'$. Lossless approaches on the other hand ensure that their result set M'' is exactly the same as M' . In this paper, we present a lossless approach. To the best of our knowledge, only one other link discovery framework implements a lossless approach that has been designed to exploit the bound defined by the threshold θ to ensure a more efficient computation of the Jaro-Winkler distance, i.e., the SILK framework with the approach MultiBlock [10]. We thus compare our approach with SILK 2.6.0 in the evaluation section of this paper.

2.2. The Jaro-Winkler Similarity

Let Σ^* be the set of all possible strings over Σ . The Jaro measure $d_j : \Sigma^* \times \Sigma^* \rightarrow [0, 1]$ is a string similarity measure approach which was developed originally for name comparison in the U.S. Census. This measure takes into account the number of character matches m and the ratio of their transpositions t :

$$d_j = \begin{cases} 0 & \text{if } m = 0 \\ \frac{1}{3} \left(\frac{m}{|s|} + \frac{m}{|t|} + \frac{m-\tau}{m} \right) & \text{otherwise} \end{cases} \quad (1)$$

Here two characters are considered to be a match if and only if (1) they are the same and (2) they are at most at a distance $w = \lfloor \frac{\max(|s|, |t|)}{2} \rfloor$ from each other. For example, for $s = \text{spears}$ and $t = \text{pears}$, the second s of s matches the s of t while the first s of s does not match the s of t . The number of transpositions τ is half of the number of matching characters which are in different orders across the input strings s and t . For ex-

⁸<http://limes.sf.net>

ample, consider strings `house` and `huose`. All characters match, however, the order of the letters `o` and `u` is different. Thus, we have $\tau = \frac{2}{2} = 1$.

The Jaro-Winkler measure [30] is an extension of the Jaro distance. This extension is based on Winkler's observation that typing errors occur most commonly in the middle or at the end of a word, but very rarely in the beginning. Hence, it is legitimate to put more emphasis on matching prefixes if the Jaro distance exceeds a certain "boost threshold" b_t , originally set to 0.7 by Winkler in his publication [30].

$$d_w = \begin{cases} d_j & \text{if } d_j < b_t \\ d_j + (\ell p (1 - d_j)) & \text{otherwise} \end{cases} \quad (2)$$

Here, ℓ denotes the length of the common prefix and p is a weighting factor. Winkler uses $p = 0.1$ and $\ell \leq 4$. Note that ℓp must not be greater than 1.

3. Improving the Runtime of Bounded Jaro-Winkler

The main principle behind reducing the runtime of the computation of measures is to improve their reduction ratio. Here, we use a sequence of filters that allow discarding similarity computations while being sure that they would have led to a similarity score which would have been less than our threshold θ . To this end, we regard the problem as that of finding filters that return an upper bound estimate $\theta_e(s, t) \geq d_w(s, t)$ for some properties of the input strings that can be computed in constant time. For a given threshold θ , if $\theta_e(s, t) \leq \theta$, we can discard the computation of the Jaro-Winkler score for (s, t) . We denote the upper bound estimate of a concrete filter by $\theta_{e,x}$, where x will be the first letter of the filters name.

3.1. Length-based filters

In the following, we denoted the length of a string s with $|s|$. Our first filter is based on the insight that large length differences are a guarantee for poor similarity. For example, the strings `a` and `alpha` cannot have a Jaro-Winkler similarity of 1 by virtue of their length difference. We can formalize this idea as follows: Let s and t be strings with respective lengths $|s|$ and $|t|$. Without loss of generality, we will assume that $|s| \leq |t|$. Moreover, let m be the number of matches across s and t . Because $m \leq |s|$, we can substitute m with

$|s|$ and gain the following upper bound estimation for $d_j(s, t)$:

$$\begin{aligned} d_j &= \frac{1}{3} \left(\frac{m}{|s|} + \frac{m}{|t|} + \frac{m - \tau}{m} \right) \\ &\leq \frac{1}{3} \left(1 + \frac{|s|}{|t|} + \frac{|s| - \tau}{|s|} \right) \end{aligned} \quad (3)$$

Now the lower bound for the number τ of transpositions is 0. Thus, we obtain the following equation.

$$d_j \leq \frac{1}{3} \left(1 + \frac{|s|}{|t|} + 1 \right) \leq \frac{2}{3} + \frac{|s|}{3|t|} \quad (4)$$

The application of this approximation on Winkler's extension is trivial:

$$\begin{aligned} d_w &= d_j + \ell \cdot p \cdot (1 - d_j) \\ &\leq \frac{2}{3} + \frac{|s|}{3|t|} + \ell \cdot p \cdot \left(\frac{1}{3} - \frac{|s|}{3|t|} \right) = \theta_{e,l} \end{aligned} \quad (5)$$

Consider the pair $s = \text{bike}$ and $t = \text{bicycle}$ and a threshold $\theta = 0.9$. Applying the estimation for Jaro we get $d_j \leq \frac{2}{3} - \frac{4}{3 \cdot 7} = 0.857$. This exceeds the boost threshold, so we use equation 5 to compute $\theta_{e,l}(s, t) = 0.885$. Given that $\theta_{e,l}(s, t) < \theta$, we do not need to compute the real value of $d_w(s, t)$ as we know that this pair will not belong to the desired output.

By using this approach we can decide in $O(1)$ ⁹ if a given pairs score is greater than a given threshold, which saves us the much more expensive score computation for a big number of pairs, provided that the input strings sufficiently vary in length.

3.2. Filtering ranges by length

The approach described above can be reversed to limit the number of pairs that we are going to be iterating over. To this end, we can construct an *index* : $\mathbb{N} \rightarrow \Sigma^*$ which maps string lengths $l \in \mathbb{N}$ to all strings s with $|s| = l$. With the help of this index, we can now determine the set of strings t that should be compared with the subset $S(l)$ of S that only contains strings of length l . We go about using this insight by computing the upper and lower bound for the length of a string t

⁹In most programming languages, especially Java (which we used for our implementation), the length of string is stored in a variable and can thus be accessed in constant time.

that should be compared with a string s . This is basically equivalent to asking what is the minimum length difference $||s| - |t||$ so that $\theta \geq \theta_{e,l}(s, t)$ is satisfied. We transpose equation 5 to the following for our lower bound:

$$|s| \geq \left\lceil 3|t| \frac{\theta - \ell p}{1 - \ell p} - 2|t| \right\rceil = \rho_{min} \quad (6)$$

Analogously, we can derive the following upper bound:

$$|s| \leq \left\lceil \frac{|t|}{3 \frac{\theta - \ell p}{1 - \ell p} - 2} \right\rceil = \rho_{max} \quad (7)$$

For example, consider a list of strings S with equally distributed string lengths (4, 7, 11, 18) and $\theta = 0.9$. Using Equation 6 and Equation 7 we obtain Table 1. Taking into account the last column of the table, we

Table 1
Bounds for distinct string lengths ($\theta = 0.9$)

$ s $	ρ_{min}	ρ_{max}	sizes in range
4	2	8	(4, 7)
7	3	14	(4, 7, 11)
11	5	22	(7, 11, 18)
18	9	36	(11, 18)

will save a total of $\frac{3}{8}$ comparisons. Note that this bound is not effective if the threshold is not greater than $\phi = \frac{2+\ell p}{3}$. This is simply due to the bounds ρ_{min} and ρ_{max} being negative below this particular value of ϕ .

3.3. Filtering by character frequency

An even more fine-grained approach can be chosen to filter out computations. Let $e : \Sigma \times \Sigma^* \rightarrow \mathbb{N}$ be the function which returns the number of occurrences of a given character in a string. For the strings s and t , the number of maximum possible matches m_{max} can be expressed as

$$m_{max} = \sum_{c \in S} \min(e(s, c), e(t, c)) \geq m \quad (8)$$

Consequently, we can now substitute m for m_{max} in the Jaro distance computation:

$$\begin{aligned} d_j(s, t) &= \frac{1}{3} \left(\frac{m_{max}}{|s|} + \frac{m_{max}}{|t|} + \frac{m_{max} - \tau}{m_{max}} \right) \\ &\leq \frac{1}{3} \left(\frac{m_{max}}{|s|} + \frac{m_{max}}{|t|} + 1 \right) = \theta_{e,f} \end{aligned} \quad (9)$$

Table 2

Calculation of m_{max}

c	$e(s, c)$	$e(t, c)$	$\min(e(s, c), e(t, c))$	m_{max}
a	2	1	1	1
c	0	1	0	1
h	0	1	0	1
i	0	1	0	1
m	0	1	0	1
n	1	0	0	1
o	1	1	1	2
p	0	1	0	2
r	1	1	1	3
s	1	1	1	4
t	2	1	1	5
u	1	0	0	5

We can thus derive that $d_j(s, t) \geq \theta$ iff

$$m_{max} \geq \frac{(3\theta - 1) |s||t|}{|s| + |t|}. \quad (10)$$

For instance, let $s = \text{astronaut}$, $t = \text{astrochimp}$. The retrieval of m_{max} is shown in Table 2.

3.3.1. Naïve Implementation

The naïve implementation of the character-based filter consists of checking if Equation 10 holds true. The e function for each string is thereby represented using a map. As shown by our evaluation, the character-based filter leads to a significant reduction of the number of comparisons (see Figure 6) by more than 2 orders of magnitude. However, the runtime improvement achieved using this implementation is not substantial. This is simply due to the lookup into maps being constant in time complexity but still a large amount of time. Instead of regarding strings as monolithic entities, we thus extended our implementation to be more fine-grained and used a trie as explained in the subsequent section.

3.3.2. Implementation with Tries

To overcome the need to perform character index lookups for every pair of strings we use a trie-based pruning technique. We thus dub this filter the *trie filter* from now on. Let $\mathcal{D}_l \in \{S, T\}$ be the dataset that contains the longest string. Moreover, let \mathcal{D}_s be the other dataset (we set $D_l = S$ and $D_s = T$ if both datasets contain longest strings of same length).

We define a function $\sigma : \Sigma^+ \rightarrow \Sigma^+$, which maps a word onto an ordered permutation of itself according to any consistent total ordering of letters in Σ . For ex-

Algorithm 1 Trie filter

```

1: procedure TRIEFILTER( $\mathcal{D}_l, \mathcal{D}_s, \theta$ )
2:    $\mathcal{M} \leftarrow \{\}$ 
3:    $\mathcal{T} \leftarrow \text{CONSTRUCTTRIE}(\mathcal{D}_s)$ 
4:    $\mathcal{P} \leftarrow \text{CONSTRUCTPARTITIONS}(\mathcal{D}_l)$ 
5:   for all  $\mathcal{P}_u \in \mathcal{P}$  do
6:     TRIETRAVERSE( $\mathcal{T}, \mathcal{P}_u, u, \theta$ )
7:   end for
8: end procedure

9: function  $\theta_e, t(l_1, l_2, m)$ 
10:   $p \leftarrow 4$ 
11:   $\ell \leftarrow 0.1$ 
12:  return  $\frac{1-\ell p}{3} \left( \frac{m}{l_1} + \frac{m}{l_2} + 1 \right) + \ell p$ 
13: end function

14: procedure TRIETRAVERSE( $\mathcal{T}, \mathcal{P}_u, u, \theta$ )
15:   $\mathcal{X} \leftarrow \text{new Stack}()$ 
16:   $\mathcal{X}.\text{push}(\langle \epsilon, 0, 0 \rangle)$   $\triangleright$  Start at root node
17:  while  $\mathcal{X} \neq \emptyset$  do
18:     $\langle \eta, m_u, d_u \rangle \leftarrow \mathcal{X}.\text{pop}()$ 
19:     $m_{max} \leftarrow m_u + \min(|u| - d_u, \text{height}(\eta) + 1)$ 
20:    if  $m_{max} \geq m_{min}(|u|, v)$  then
21:      if  $\text{label}(\eta) \leq u_{d_u}$  then
22:        if  $\text{label}(\eta) = u_{d_u}$  then
23:           $d_u \leftarrow d_u + 1$ 
24:           $m_u \leftarrow m_u + 1$ 
25:           $\alpha_1 \leftarrow \mathcal{S}_\eta \neq \emptyset$ 
26:           $\alpha_2 \leftarrow d_w^{max}(|\eta|, |u|, m) \geq \theta$ 
27:             $\triangleright |\eta|$  denotes the level of  $\eta$ 
28:          if  $\alpha_1 \wedge \alpha_2$  then
29:             $\mathcal{M} \leftarrow \mathcal{M} \cup \langle \mathcal{S}_\eta, \mathcal{P}_u \rangle$ 
30:          end if
31:        end if
32:        for all  $x \in \text{children}(\eta)$  do
33:           $\mathcal{X}.\text{push}(\langle x, m_u, d_u \rangle)$ 
34:        end for
35:      else
36:         $d_u \leftarrow d_u + 1$ 
37:         $\mathcal{X}.\text{push}(\langle \eta, m_u, d_u \rangle)$ 
38:      end if
39:    end while
40: end procedure

```

ample, we get $\sigma(\text{hello}) = \text{ehlllo}$ for the ordering derived from the UTF-8 index of characters. We begin by partitioning the set \mathcal{D}_l . The partition (also called bucket) $\mathcal{P}_{\sigma(s_i)}$ with the signature $\sigma(s_i)$ is the set of all strings s_j from \mathcal{D}_l with $\sigma(s_i) = \sigma(s_j)$. In a second step, we add all elements $t_k \in \mathcal{D}_s$ to a trie \mathcal{T} . More precisely, we add all t_k to sets \mathcal{S}_η associated with nodes η at path $\sigma(t_k)$. We assume that each node of \mathcal{T} is labelled with the letter that must be added to its father to create it. The core of our approach now consists of detecting all t_k that must be compared with all elements of a particular partition $\mathcal{P}_{\sigma(s_i)}$. To this end, we begin by computing the minimal number of matching characters m_{min} from the signature $\sigma(s_i)$ that any relevant t_k must have to be a potential match. The total amount of possible matches m_{max} must be greater than or equal to m_{min} for having a chance of finding any t_k such that $d_w(s_i, t_k) \geq \theta$.

From the Jaro-Winkler version of the upper bound estimation given in Equation 9 we obtain the following

equation:

$$\begin{aligned}
 m_{max} &\geq m_{min}(s, t) \\
 &= \left\lceil \left(\theta - \ell p - \frac{1 - \ell p}{3} \right) \cdot \frac{3|s||t|}{(|s| + |t|)(1 - \ell p)} \right\rceil
 \end{aligned}
 \tag{11}$$

Given this value, we traverse \mathcal{T} depth-first in pre-order for every \mathcal{P}_u with signature u (i.e., for all strings s_i with signature u) as follows: We begin by setting the number m_u of matches for u to 0 and the depth counter d_u to 0. Note that d_u and m_u do have a local scope w.r.t. a path on \mathcal{T} . That is, we virtually run our approach independently on each path. When at a node η of the trie, we check whether the number of matches between the sequence of letters that leads to η and the signature u abides by $m_{max} \geq m_{min}$. Less formally, we simply check whether the children of the given nodes can at all be potential matches for the signature u . If the node η does not fulfill this criterion, we terminate the exploration of the trie and go on with the next node in the set of nodes to be explored, effectively cutting off the subtree below node η . Else, we make use

of the total ordering underlying the signature u . To this end, we compare the letter of u with the index d_u , from here on called u_{d_u} with the label of η . If the label of η is superior to that chosen letter u_{d_u} , then no children of η can match u_{d_u} either. However, subsequent letters of u may match the label of η or its children. We therefore increment d_u and put η back into the set of nodes to be explored, effectively revisiting η in the next iteration step. If the label of η is inferior to u_{d_u} , it means that we can still expect some node in the subtree of η to match u_{d_u} , therefore we just add all children of η to the set of nodes to be explored. Finally, if the label of η matches u_{d_u} exactly, we increment both m_u and d_u , add all children of η to the set of nodes to be explored and if \mathcal{S}_η is not empty, we add $(\mathcal{P}_u, \mathcal{S}_\eta)$ to the set of potential match candidates. The detailed implementation logic of the tree traversal is given in Algorithm 1. Consider the following example for a better understanding of how the algorithm works. The threshold parameter θ is set to 0.92. Datasets S and T are given in Table 3.

Table 3
Datasets S and T from the example

S	T
{nines}	{niece, niche, nices, spice, since}

We deliberately selected string of size 5 so that we can present the character filtering in a clear fashion. As both datasets contain the longest string, we set $D_l = S$ and $D_s = T$. The only string in D_l is `nines`, hence we have exactly one Partition $\mathcal{P}_{\sigma(\text{nines})} = \mathcal{P}_{\text{einns}} = \{\text{nines}\}$. We now construct \mathcal{T} from strings in D_s , gaining the trie shown in Figure 1. The nodes η , whose corresponding sets \mathcal{S}_η are not empty, are highlighted in blue color. Table 4 gives an overview of all \mathcal{S}_η .




Fig. 1. Example Trie \mathcal{T}

Table 4
Nodes and their corresponding sets from the example

η	\mathcal{S}_η
s_0	{spice}
s_1	{nices, since}
n_0	{niche}
n_1	{niece}

We select the first and only Partition $\mathcal{P}_{\text{einns}}$ and set $u = \text{einns}$, $d_u = 0$ and $m_u = 0$. Now we begin the trie exploration, therefore starting at the root (labeled with the empty word). In the next step we need to check whether $m_{max} \geq m_{min}$ holds true. To this end, we need to compute m_{min} and m_{max} . First, we compute m_{min} using Equation 11 with arguments u and v , where v is the string we gain by chaining all labels of nodes on the shortest path from the root to a leaf, getting $m_{min} = 4$. Second, we determine the maximal number of possible matches in this iteration using Equation 12.

$$m_{max}(u, d_u, \eta) = m + \min(|u| - d_u, \text{height}(\eta) + 1) \quad (12)$$

The height of a node η herein is defined as the number of edges on the longest downward path between that node and a leaf. Therefore we have $m_{max} = \min(5, 6) = 6$. As $m_{max} = 6 \geq 4 = m_{min}$ holds true, we continue by comparing the label of η to u_{d_u} . With η being the root of the trie, we have $\text{label}(\eta) = \epsilon < e = u_0$, hence we just add all children of η to the set of nodes to be explored.

In the next iteration step, we now have $m_{max} = \min(5, 5) = 5$, which is still greater than $m_{min} = 4$. Again, we have $\text{label}(\eta) = c < e = u_0$ and end up adding all children of node c to the set of nodes to be explored.

In iteration step 3 $m_{max} = 4 \geq 4 = m_{min}$ still holds true and this time we indeed have a match, as $\text{label}(\eta) = e = u_0$. Therefore, we now increment both m and d_u and add all children of node e_0 to the set of nodes to be explored.

In iteration step 4 we select $\eta = i_0$ and still $m_{max} = 4 \geq 4 = m_{min}$. We have a match for $\text{label}(\eta) = i = u_1$. We increment both m and d_u and add all children of node i_0 to the set of nodes to be explored.

In iteration step 5 we select $\eta = p$ and still $m_{max} = 4 \geq 4 = m_{min}$. We have $\text{label}(\eta) = p > n = u_2$. We increment d_u and revisit this node in the next step.




Fig. 2. Flowchart of parallel trie filter stack

In iteration step 6 we revisit $\eta = p$ and still $m_{max} = 4 \geq 4 = m_{min}$. Again, it is $label(\eta) = p > n = u_3$. We increment d_u and revisit this node once again in the next step.

In iteration step 7 we revisit $\eta = p$ and this time $m_{max} = 3 \geq 4 = m_{min}$, hence we continue with the next node.

In iteration step 8 we select $\eta = n_0$ and still $m_{max} = 4 \geq 4 = m_{min}$. We have a match for $label(\eta) = n = u_2$. Therefore, we now increment both m and d_u and add all children of node n_0 to the set of nodes to be explored.

In iteration step 9 we select $\eta = s_1$ and still $m_{max} = 4 \geq 4 = m_{min}$. We have $label(\eta) = s > n = u_3$. We increment d_u and revisit this node in the next step.

In iteration step 10 we revisit $\eta = s_1$ and still $m_{max} = 4 \geq 4 = m_{min}$. We have a match for $label(\eta) = s = u_4$. As $m_u = 4 = m_{min}$, we can now add $\{\mathcal{P}_u, \mathcal{S}_{s_1}\}$ to the set of match candidates.

In iteration step 11 we select $\eta = h$ and still $m_{max} = 4 \geq 4 = m_{min}$. However, $label(\eta) = h < i = u_1$. All children of node h are therefore added to the set of nodes to be explored.

In iteration step 12 we select $\eta = i_1$ and due to $m_{max} = 3 < 4 = m_{min}$, we abort the exploration of this branch and continue with the next step.

Iterations 13 and 14 come about analogously to steps 11 and 12, leading to a stop of exploration at node i_2 . The set of match candidates that is forwarded to the actual Jaro-Winkler distance computation now contains $\{\mathcal{P}_u, \mathcal{S}_{s_1}\} = \{\{nines\}, \{nices, since\}\}$.

Note that using the range filter prior to the trie filter is a key requirement towards its efficiency, that is, the more equally distributed the string lengths are, the better we get in terms of reduction ratio and runtime improvement. The worst case scenario is a big dataset with very large strings over a little alphabet and little to none variation in string lengths.

3.4. Parallel Implementation



While our results suggest that the approach presented above scales well on one processor, we wanted to measure how well the partitioning induced by the upper and lower length bounds given in Equations 6 and 7 can be used to implement our approach in parallel. Let $T_{i,j} = \{t \in T : i \leq |t| \leq j\}$ be set of strings whose length is larger or equal to i and less or equal to j . Moreover, let $S_k = \{s \in S : k = |s|\}$ be the subset of S which contains strings of length k . Finally, let $\mathcal{L}(S)$ be the set of distinct string lengths of elements of S . We distribute a subset $S \times T$ into sets $M_x''' = S_x \times T_{\rho_{min}(x), \rho_{max}(x)} \forall x \in \mathcal{L}(S)$. Note that due to the bounds in Equations 6 and 7, we are sure that all elements of T which abide by the similarity condition $\delta(s, t) \geq \theta$ for $s \in S_x$ can be found in the subset $T_{\rho_{min}(x), \rho_{max}(x)}$. Moreover, not all elements of $S \times T$ are in the set M_x''' . In particular, pairs of strings that do not abide by the range restrictions are not considered as they are known not to be matches.

Now given the sets M_x''' , we can parallelize our implementation by simply using a thread-pool-based approach. We initialize a user-given number of threads and assign each thread one of the sets M_x''' randomly. Once a thread has completed its computations, it is simply assigned a further set M_x''' . Note that we do not implement any load balancing as we were primarily interested in how well the range filter allows partitioning data. Figure 2 gives an overview of our parallel pipeline.

4. Evaluation

4.1. Experimental Setup

The aim of our evaluation was to study how well our approach performs on real data. We chose DBpedia 3.9 as a source of data for our experiments as it contains data pertaining to 1.1 million persons and thus

Fig. 3. Runtimes on sample of DBpedia `rdf:labels` with size 1000Fig. 4. Runtimes on sample of DBpedia `rdf:labels` with size 10^4

allows for both fine-grained evaluations and scalability evaluations. As a second data source we chose LinkedGeoData to evaluate how well we perform on strings that have no relation to person names. We chose these datasets because (1) they have been widely used in experiments pertaining to link discovery and (2) the distributions of string sizes in these datasets are significantly different (see Figures 7 and 8). All experiments were deduplication experiments, i.e., $S = T$. We considered the list of all `rdfs:label` in DBpedia in our runtime evaluation and added all `rdfs:labels` of the Places dataset from LinkedGeoData for scalability experiments. We also computed the number of actual Jaro-Winkler calculations carried out for 1000 strings from DBpedia. All runtime and scalability ex-

periments were performed on a 2.5 GHz Intel Core i5 machine with 16GB RAM running OS X 10.9.3. The speedup of the parallel trie-based filter stack was measured on a Microsoft Azure VM instance with 16 CPUs and 112GB RAM.

For the sake of legibility, we used the following names and acronyms for filter combination or lack thereof:

- naïve (`n`): No use of any filters (equivalent to brute-force).
- range (`r`): Uses only the range filter as described in Section 3.2.
- length (`l`): Uses only the length filter as described in Section 3.1.

- frequency (f): Uses only the naïve implementation of the character frequency filter as described in Section 3.3.
- trie (t): Uses only the trie filter implementation of the character frequency filter as described in Section 3.3.2.
- range and length (r+l): Uses the range filter to partition the datasets, the length filter is then applied to each pair of partitions.
- range and frequency (r+f): Sets up the frequency filter globally in a first step. Then uses the range filter to partition the datasets and applies the frequency filter to each pair of partitions.
- range and trie (r+t): Uses the range filter to partition the datasets, the trie filter is then applied to each pair of partitions.
- range, length and frequency (r+l+f): Sets up the frequency filter globally in a first step. Then uses the range filter to partition the datasets and applies the length filter to each pair of partitions. The frequency filter is then only applied to those pairs of strings for which the length filter does not apply.

Note that the trie filter implicitly contains a length filter, which is the reason why we do not evaluate the combination $r+l+t$ separately.

4.2. Runtime Evaluation

In our first series of experiments, we evaluated the runtime of all filter combinations against the naïve approach on a small dataset containing 1000 labels from DBpedia. The results of our evaluation are shown in Figure 3. This evaluation suggests that all filter setups except those containing t outperform the naïve approach. Moreover, the combination of all filters leads to the best overall runtime in most cases on this small sample. Overall, the results on this dataset already shows that we outperform the naïve approach by more than an order of magnitude when $\theta > 0.9$. Interestingly, the break-even point for f and t is reached when $\theta > 0.99$ on this small dataset. This is clearly due to the overhead necessary to create the trie overshadowing the runtime advantage engendered by using the trie to search for matches.

The runtimes on a larger sample of size 10^4 show an even better improvement (see Figure 4). This suggests that the relative improvement of our approach improves with the size of the problem. The most interesting result comes from filter setup $r+t$. It is slower

than the naïve approach on low θ but after a break-even point around $0.89 \leq \theta \leq 0.91$ it outperforms the second best setup $r+l$ by an order of magnitude.

4.3. Scalability Evaluation

The aim of the scalability evaluation was to measure how well our approach scales. In our first set of experiments, we looked at the growth of the runtime of our approach on datasets of growing sizes (see Figures 9 and 10). Our results show very clearly that $r+t$ is the best filter combination for datasets of large sizes. This result holds on both DBpedia and LinkedGeoData. $r+t$ is thus the default implementation of the Jaro-Winkler measure in LIMES. In addition, our results suggest that our approach grows linearly with the number of labels contained in S and T . This is one of our most important results as it makes clear that we can employ $r+t$ on large datasets and expect acceptable runtimes. Moreover, the behavior of the runtime of our default setup can be easily predicted for large datasets, which is of importance when asking users to wait for the results of the computation.

The second series of scalability experiments looked at the runtime behaviour of our approach on a large dataset with 10^5 labels (see Figure 5). Our results suggest that the runtime of our approach falls superlinearly with an increase of the threshold θ . This behaviour suggest that our approach is especially useful on clean datasets, where high thresholds can be used for link discovery.

In the third series of experiments we looked at the speedup we gain by parallelizing $r+t$ on the DBpedia dataset with input sizes of 10^5 and 10^6 (see Figures 11 and 12). Here, our results show that the current implementation scales up in a satisfactory manner on up to 4 processors. Running the parallel implementation on 8 and 12 processors does not bring about any considerable increase in speedup. The reason for this behavior is simply that we did not implement any load balancing. Hence, there is commonly one thread that is assigned a single large M_x''' , leading to all other threads having completed their tasks but having to wait for this particular thread to terminate. Adding load balancing to the approach as well as splitting large M_x''' into smaller chunks should lead to an improvement of the scalability of our parallel implementation. These extensions will be implemented in future work.




Fig. 5. Runtimes on sample of DBpedia `rdf:labels` with size 10^5 , scaling threshold




Fig. 6. Actual Jaro-Winkler computations, measured on sample of DBpedia `rdf:labels` with size 1000




Fig. 7. String length distribution of 10^6 DBpedia `rdf:labels`




Fig. 8. String length distribution of 10^6 LinkedGeoData `rdf:labels`

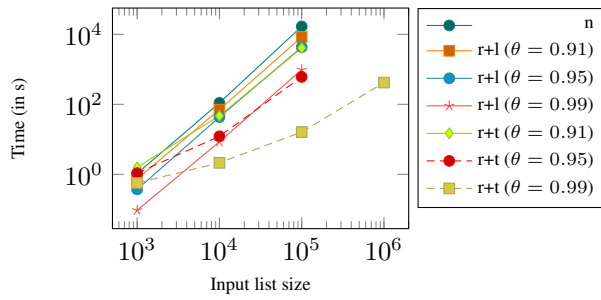


Fig. 9. Runtimes for growing input list sizes (DBPedia)

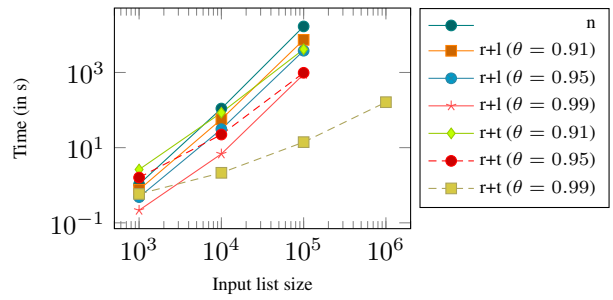


Fig. 10. Runtimes for growing input list sizes (LinkedGeoData)




Fig. 11. Speedup of parallel algorithm (DBPedia, 10^5)




Fig. 12. Speedup of parallel algorithm (DBPedia, 10^6)

4.4. Comparison with existing approaches

We compared our approach with SILK2.6.0. To this end, we retrieved all `rdfs:label` of instances of subclasses of `Person`. We only compared with SILK on small datasets (i.e., on classes with small numbers of instances) as the results on these small datasets already showed that we outperform SILK consistently.¹⁰ Our results are shown in Table 5. They suggest that the absolute difference in runtime grows with the size of the datasets. Thus, we did not consider testing larger datasets against SILK as in the best case, we were already 4.7 times faster than SILK (Architect dataset, $\theta = 0.95$).

5. Related Work

The work presented herein is related to record linkage, deduplication, link discovery and the efficient computation of Hausdorff distances. An extensive amount of literature has been published by the database community on record linkage (see [12,7] for surveys). With regard to *time complexity*, time-efficient deduplication algorithms such as PPJoin+ [32], EDJoin [31], PassJoin [13] and TrieJoin [29] were developed over the last years. Several of these were then integrated into the hybrid link discovery framework LIMES [17]. Moreover, dedicated time-efficient approaches were developed for LD. For example, RDF-AI [27] implements a five-step approach that comprises the preprocessing, matching, fusion, interlink and post-processing of data sets. [19] presents an approach based on the Cauchy-Schwarz that allows discarding a large number of unnecessary computations. The approaches HYPPO [15] and \mathcal{HR}^3 [16] rely on space tiling in spaces with measures that can be split into independent measures across the dimensions of the problem at hand. Especially, \mathcal{HR}^3 was shown to be the first approach that can achieve a relative reduction ratio r' less or equal to any given relative reduction ratio $r > 1$. Standard blocking approaches were implemented in the first versions of SILK and later replaced with MultiBlock [10], a lossless multi-dimensional blocking technique. KnoFuss [22] also implements blocking techniques to achieve acceptable runtimes. Further approaches can be found in [28,4,23,24,8,26].

In addition to addressing the runtime of link discovery, several machine-learning approaches have been

developed to learn link specifications (also called linkage rules) for link discovery. For example, machine-learning frameworks such as FEBRL [2] and MARLIN [1] rely on models such as Support Vector machines [3] and decision trees [25] to detect classifiers for record linkage. RAVEN [20] relies on active learning to detect linear or Boolean classifiers. The EAGLE approach [21] combines active learning and genetic programming to detect link specifications. KnoFuss [22] goes a step further and presents an unsupervised approach based on genetic programming for finding accurate link specifications. Other record deduplication approaches based on active learning and genetic programming are presented in [5,9].

6. Conclusion and Future Work

In this paper, we present a novel approach for the efficient execution of bounded Jaro-Winkler computations. Our approach is based on three filters which allow discarding a large number of comparisons. We showed that our approach scales well with the amount of data it is faced with. Moreover, we showed that our approach can make effective use of large thresholds by reducing the total runtime of the approach considerably. We also compared our approach with the state-of-the-art framework SILK 2.6.0 and showed that we outperform it on all datasets. In future work, we will test whether our approach improves the accuracy of specification detection algorithms such as EAGLE. Moreover, we will focus on improving the quality of matches. To this end we will split input strings into tokens and use a hybrid approach as proposed by Monge and Elkan [14], which adds to complexity of the algorithm, hence allowing for further runtime improvements. Furthermore, we will extend the parallel implementation with load balancing.

References

- [1] Mikhail Bilenko and Raymond J. Mooney. Adaptive duplicate detection using learnable string similarity measures. In Lise Getoor, Ted E. Senator, Pedro M. Domingos, and Christos Faloutsos, editors, *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 24 - 27, 2003*, pages 39–48. ACM, 2003. doi:10.1145/956750.956759.
- [2] Peter Christen. Febrl -: an open source data cleaning, deduplication and record linkage system with a graphical user interface. In Ying Li, Bing Liu, and Sunita Sarawagi, editors, *Proceedings of the 14th ACM SIGKDD International Confer-*

¹⁰We ran SILK with `-Dthreads = 1` for the sake of fairness.

Table 5
Runtimes (in seconds) of our approach (OA) and SILK 2.6.0

DBpedia Class	Size	OA(0.8)	OA(0.9)	OA(0.95)	SILK(0.8)	SILK(0.9)	SILK(0.95)
Actors	9509	15.07	10.13	6.38	27	25	25
Architect	3544	5.58	5.48	2.32	11	11	11
Criminal	5291	11.54	7.77	4.52	18	18	18

- ence on Knowledge Discovery and Data Mining, Las Vegas, Nevada, USA, August 24-27, 2008, pages 1065–1068. ACM, 2008. doi:10.1145/1401890.1402020.
- [3] Nello Cristianini and Elisa Ricci. Support vector machines. In Ming-Yang Kao, editor, *Encyclopedia of Algorithms*. Springer, 2008. doi:10.1007/978-0-387-30162-4_415.
- [4] Philippe Cudré-Mauroux, Parisa Haghani, Michael Jost, Karl Aberer, and Hermann de Meer. idmesh: graph-based disambiguation of linked data. In Juan Quemada, Gonzalo León, Yoëlle S. Maarek, and Wolfgang Nejdl, editors, *Proceedings of the 18th International Conference on World Wide Web, WWW 2009, Madrid, Spain, April 20-24, 2009*, pages 591–600. ACM, 2009. doi:10.1145/1526709.1526789.
- [5] Junio de Freitas, Gisele L. Pappa, Altigran Soares da Silva, Marcos André Gonçalves, Edleno Silva de Moura, Adriano Veloso, Alberto H. F. Laender, and Moisés G. de Carvalho. Active learning genetic programming for record deduplication. In *Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2010, Barcelona, Spain, 18-23 July 2010*, pages 1–8. IEEE, 2010. doi:10.1109/CEC.2010.5586104.
- [6] Kevin Dreßler and Axel-Cyrille Ngonga Ngomo. Time-efficient execution of bounded jaro-winkler distances. In Pavel Shvaiko, Jérôme Euzenat, Ming Mao, Ernesto Jiménez-Ruiz, Juanzi Li, and Axel Ngonga, editors, *Proceedings of the 9th International Workshop on Ontology Matching collocated with the 13th International Semantic Web Conference (ISWC 2014), Riva del Garda, Trentino, Italy, October 20, 2014.*, volume 1317 of *CEUR Workshop Proceedings*, pages 37–48. CEUR-WS.org, 2014.
- [7] Ahmed K. Elmagarmid, Panagiotis G. Ipeirotis, and Vasilios S. Verykios. Duplicate record detection: A survey. *IEEE Trans. Knowl. Data Eng.*, 19(1):1–16, 2007. doi:10.1109/TKDE.2007.250581.
- [8] Jérôme Euzenat, Alfio Ferrara, Willem Robert van Hage, Laura Hollink, Christian Meilicke, Andriy Nikolov, Dominique Ritzke, François Scharffe, Pavel Shvaiko, Heiner Stuckenschmidt, Ondrej Sváb-Zamazal, and Cássia Trojahn dos Santos. Results of the ontology alignment evaluation initiative 2011. In Pavel Shvaiko, Jérôme Euzenat, Tom Heath, Christoph Quix, Ming Mao, and Isabel F. Cruz, editors, *Proceedings of the 6th International Workshop on Ontology Matching, Bonn, Germany, October 24, 2011*, volume 814 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2011.
- [9] Robert Isele and Christian Bizer. Learning expressive linkage rules using genetic programming. *PVLDB*, 5(11):1638–1649, 2012.
- [10] Robert Isele, Anja Jentzsch, and Christian Bizer. Efficient multidimensional blocking for link discovery without losing recall. In Amélie Marian and Vasilis Vassalos, editors, *Proceedings of the 14th International Workshop on the Web and Databases 2011, WebDB 2011, Athens, Greece, June 12, 2011*, 2011.
- [11] Matthew A Jaro. Advances in record-linkage methodology as applied to matching the 1985 census of tampa, florida. *Journal of the American Statistical Association*, 84(406):414–420, 1989.
- [12] Hanna Köpcke and Erhard Rahm. Frameworks for entity matching: A comparison. *Data Knowl. Eng.*, 69(2):197–210, 2010. doi:10.1016/j.datak.2009.10.003.
- [13] Guoliang Li, Dong Deng, Jiannan Wang, and Jianhua Feng. PASS-JOIN: A partition-based method for similarity joins. *PVLDB*, 5(3):253–264, 2011.
- [14] Alvaro E. Monge and Charles Elkan. An efficient domain-independent algorithm for detecting approximately duplicate database records. In *DMKD*, page 0, 1997.
- [15] Axel-Cyrille Ngonga Ngomo. A time-efficient hybrid approach to link discovery. In Pavel Shvaiko, Jérôme Euzenat, Tom Heath, Christoph Quix, Ming Mao, and Isabel F. Cruz, editors, *Proceedings of the 6th International Workshop on Ontology Matching, Bonn, Germany, October 24, 2011*, volume 814 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2011.
- [16] Axel-Cyrille Ngonga Ngomo. Link discovery with guaranteed reduction ratio in affine spaces with minkowski measures. In Philippe Cudré-Mauroux, Jeff Heflin, Evren Sirin, Tania Tudorache, Jérôme Euzenat, Manfred Hauswirth, Josiane Xavier Parreira, Jim Hendler, Guus Schreiber, Abraham Bernstein, and Eva Blomqvist, editors, *The Semantic Web - ISWC 2012 - 11th International Semantic Web Conference, Boston, MA, USA, November 11-15, 2012, Proceedings, Part I*, volume 7649 of *Lecture Notes in Computer Science*, pages 378–393. Springer, 2012. doi:10.1007/978-3-642-35176-1_24.
- [17] Axel-Cyrille Ngonga Ngomo. On link discovery using a hybrid approach. *J. Data Semantics*, 1(4):203–217, 2012. doi:10.1007/s13740-012-0012-y.
- [18] Axel-Cyrille Ngonga Ngomo. ORCHID - reduction-ratio-optimal computation of geo-spatial distances for link discovery. In Harith Alani, Lalana Kagal, Achille Fokoue, Paul T. Groth, Chris Biemann, Josiane Xavier Parreira, Lora Aroyo, Natasha F. Noy, Chris Welty, and Krzysztof Janowicz, editors, *The Semantic Web - ISWC 2013 - 12th International Semantic Web Conference, Sydney, NSW, Australia, October 21-25, 2013, Proceedings, Part I*, volume 8218 of *Lecture Notes in Computer Science*, pages 395–410. Springer, 2013. doi:10.1007/978-3-642-41335-3_25.
- [19] Axel-Cyrille Ngonga Ngomo and Sören Auer. LIMES - A time-efficient approach for large-scale link discovery on the web of data. In Toby Walsh, editor, *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 2312–2317. IJCAI/AAAI, 2011.
- [20] Axel-Cyrille Ngonga Ngomo, Jens Lehmann, Sören Auer, and Konrad Höffner. RAVEN - active learning of link specifications. In Pavel Shvaiko, Jérôme Euzenat, Tom Heath,

- Christoph Quix, Ming Mao, and Isabel F. Cruz, editors, *Proceedings of the 6th International Workshop on Ontology Matching, Bonn, Germany, October 24, 2011*, volume 814 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2011.
- [21] Axel-Cyrille Ngonga Ngomo and Klaus Lyko. EAGLE: efficient active learning of link specifications using genetic programming. In Elena Simperl, Philipp Cimiano, Axel Polleres, Óscar Corcho, and Valentina Presutti, editors, *The Semantic Web: Research and Applications - 9th Extended Semantic Web Conference, ESWC 2012, Heraklion, Crete, Greece, May 27-31, 2012. Proceedings*, volume 7295 of *Lecture Notes in Computer Science*, pages 149–163. Springer, 2012. doi:10.1007/978-3-642-30284-8_17.
- [22] Andriy Nikolov, Mathieu d’Aquin, and Enrico Motta. Unsupervised learning of link discovery configuration. In Elena Simperl, Philipp Cimiano, Axel Polleres, Óscar Corcho, and Valentina Presutti, editors, *The Semantic Web: Research and Applications - 9th Extended Semantic Web Conference, ESWC 2012, Heraklion, Crete, Greece, May 27-31, 2012. Proceedings*, volume 7295 of *Lecture Notes in Computer Science*, pages 119–133. Springer, 2012. doi:10.1007/978-3-642-30284-8_15.
- [23] Andriy Nikolov, Victoria S. Uren, Enrico Motta, and Anne N. De Roeck. Overcoming schema heterogeneity between linked semantic repositories to improve coreference resolution. In Asunción Gómez-Pérez, Yong Yu, and Ying Ding, editors, *The Semantic Web, Fourth Asian Conference, ASWC 2009, Shanghai, China, December 6-9, 2009. Proceedings*, volume 5926 of *Lecture Notes in Computer Science*, pages 332–346. Springer, 2009. doi:10.1007/978-3-642-10871-6_23.
- [24] George Papadakis, Ekaterini Ioannou, Claudia Niederée, Themis Palpanas, and Wolfgang Nejdl. Eliminating the redundancy in blocking-based entity resolution methods. In Glen Newton, Michael J. Wright, and Lillian N. Cassel, editors, *Proceedings of the 2011 Joint International Conference on Digital Libraries, JCDL 2011, Ottawa, ON, Canada, June 13-17, 2011*, pages 85–94. ACM, 2011. doi:10.1145/1998076.1998093.
- [25] S. Rasoul Safavian and David A. Landgrebe. A survey of decision tree classifier methodology. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(3):660–674, 1991. doi:10.1109/21.97458.
- [26] Fatiha Saïs, Nathalie Pernelle, and Marie-Christine Rousset. Combining a logical and a numerical method for data reconciliation. *J. Data Semantics*, 12:66–94, 2009. doi:10.1007/978-3-642-00685-2_3.
- [27] Francois Scharffe, Yanbin Liu, and Chuguang Zhou. Rdf-ai: an architecture for rdf datasets matching, fusion and interlink. In *Proc. IJCAI 2009 workshop on Identity, reference, and knowledge representation (IR-KR), Pasadena (CA US)*, 2009.
- [28] Dezhao Song and Jeff Heflin. Automatically generating data linkages using a domain-independent candidate selection approach. In Lora Aroyo, Chris Welty, Harith Alani, Jamie Taylor, Abraham Bernstein, Lalana Kagal, Natasha Fridman Noy, and Eva Blomqvist, editors, *The Semantic Web - ISWC 2011 - 10th International Semantic Web Conference, Bonn, Germany, October 23-27, 2011, Proceedings, Part I*, volume 7031 of *Lecture Notes in Computer Science*, pages 649–664. Springer, 2011. doi:10.1007/978-3-642-25073-6_41.
- [29] Jiannan Wang, Guoliang Li, and Jianhua Feng. Trie-join: Efficient trie-based string similarity joins with edit-distance constraints. *PVLDB*, 3(1):1219–1230, 2010.
- [30] William E. Winkler. String comparator metrics and enhanced decision rules in the fellegi-sunter model of record linkage. In *Proceedings of the Section on Survey Research*, pages 354–359, 1990.
- [31] Chuan Xiao, Wei Wang, and Xuemin Lin. Ed-join: an efficient algorithm for similarity joins with edit distance constraints. *PVLDB*, 1(1):933–944, 2008.
- [32] Chuan Xiao, Wei Wang, Xuemin Lin, and Jeffrey Xu Yu. Efficient similarity joins for near duplicate detection. In Jinpeng Huai, Robin Chen, Hsiao-Wuen Hon, Yunhao Liu, Wei-Ying Ma, Andrew Tomkins, and Xiaodong Zhang, editors, *Proceedings of the 17th International Conference on World Wide Web, WWW 2008, Beijing, China, April 21-25, 2008*, pages 131–140. ACM, 2008. doi:10.1145/1367497.1367516.