

# A Comparison of Blocking Methods for Record Linkage

Rebecca C. Steorts, Samuel L. Ventura,  
Mauricio Sadinle, and Stephen E. Fienberg\*

Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213  
beka@cmu.edu, {sventura, msadinle, fienberg}@stat.cmu.edu

**Abstract.** Record linkage seeks to merge databases and to remove duplicates when unique identifiers are not available. Most approaches use blocking techniques to reduce the computational complexity associated with record linkage. We review traditional blocking techniques, which typically partition the records according to a set of field attributes, and consider two variants of a method known as locality sensitive hashing, sometimes referred to as “private blocking.” We compare these approaches in terms of their recall, reduction ratio, and computational complexity. We evaluate these methods using different synthetic datafiles and conclude with a discussion of privacy-related issues.

## 1 Introduction

A commonly encountered problem in practice is merging databases containing records collected by different sources, often via dissimilar methods. Different variants of this task are known as record linkage, de-duplication, and entity resolution. Record linkage is inherently a difficult problem [2, 11, 12]. These difficulties are partially due to the noise inherent in the data, which is often hard to accurately model [17, 20]. A more substantial obstacle, however, is the scalability of the approaches [23]. With  $d$  databases of  $n$  records each, brute-force approaches, using all-to-all comparisons, require  $O(n^d)$  comparisons. This is quickly prohibitive for even moderate  $n$  or  $d$ . To avoid this computational bottleneck, the number of comparisons made must be drastically reduced, without compromising linkage accuracy. Record linkage is made scalable by “blocking,” which involves partitioning datafiles into “blocks” of records and treating records in different blocks as non-co-referent *a priori* [2, 11]. Record linkage methods are only applied *within* blocks, reducing the comparisons to  $O(Bn_{\max}^d)$ , with  $n_{\max}$  being the size of the largest of the  $B$  blocks.

The most basic method for constructing a blocking partition picks certain fields (e.g. geography, or gender and year of birth) and places records in the same block if and only if they agree on all such fields. This amounts to an  $a$

---

\* This research was partially supported by the National Science Foundation through grants SES1130706 and DMS1043903 to the Department of Statistics, Carnegie Mellon University.

*priori* judgment that these fields are error-free. We call this *traditional blocking* (§2.1).

Other data-dependent blocking methods [2, 23] are highly application-specific or are based on placing similar records into the same block, using techniques of “locality-sensitive hashing” (LSH). LSH uses all of the information contained in each record and can be adjusted to ensure that blocks are manageably small, but then does not allow for further record linkage within blocks. For example, [15] introduced novel data structures for sorting and fast approximate nearest neighbor look-up within blocks produced by LSH. Their approach gave balance between speed and recall, but their technique is very specific to nearest neighbor search with similarity defined by the hash function. Such methods are fast and have high recall, but suffer from low precision, rather, too many false positives. This approach is called *private* if, after the blocking is performed, all candidate records pairs are compared and classified into matches/non-matches using computationally intensive “private” comparison and classification techniques [3].

Some blocking schemes involve clustering techniques to partition the records into clusters of similar records. [16] used canopies, a simple clustering approach to group similar records into overlapping subsets for record linkage. Canopies involves organizing the data into overlapping clusters/canopies using an inexpensive distance measure. Then a more expensive distance measure is used to link records within each canopy, reducing the number of required comparisons of records. [21] used a sorted nearest neighborhood clustering approach, combining  $k$ -anonymous clustering and the use of publicly available reference values to privately link records across multiple files.

Such clustering-based blocking schemes motivate our variants of LSH methods for blocking. The first, transitive locality sensitive hashing (TLSH), is based upon the community discovery literature such that a *soft transitivity* (or relaxed transitivity) can be imposed across blocks. The second,  $k$ -means locality sensitive hashing (KLSH), is based upon the information retrieval literature and clusters similar records into blocks using a vector-space representation and projections. (KLSH has been used before in information retrieval but never with record linkage [18].)

The organization of this paper is as follows. §2 reviews traditional blocking. We then review other blocking methods in §2.2 stemming from the computer science literature. §2.3 presents two different methods based upon locality sensitive hashing, TLSH and KLSH. We discuss the computational complexity of each approach in §3. We evaluate these methods (§4) on simulated data using recall, reduction ratio, and the empirical computational time as our evaluation criteria, comparing to the other methods discussed above. Finally we discuss privacy protection aspects of TLSH and KLSH, given the description of LSH as a “private” blocking technique.

## 2 Blocking Methods

Blocking divides records into mutually exclusive and jointly exhaustive “blocks,” allowing the linkage to be performed within each block. Thus, only records within the same block can be linked; linkage algorithms may still aggregate information across blocks. Traditional blocking requires domain knowledge to pick out highly reliable, if not error-free, fields for blocking. This methodology has at least two drawbacks. The first is that the resulting blocks may still be so large that linkage within them is computationally impractical. The second is that because blocks *only* consider selected fields, much time may be wasted comparing records that happen to agree on those fields but are otherwise radically different.

We first review some simple alternatives to traditional blocking on fields, and then introduce other blocking approaches that stem from computer science.

### 2.1 Simple Alternatives to Blocking

Since fields can be unreliable for many applications, blocking may miss large proportions of matches. Nevertheless, we can make use of domain-specific knowledge on the types of errors expected for field attributes. To make decisions about matches/non-matches, we must understand the *kinds of errors* that are unlikely for a certain field or a combination of them. With this information, we can identify a pair as a non-match when it has strong disagreements in a combination of fields. It is crucial that this calculation be scalable since it must be checked for all pairs of records. Some sequence of these steps reduces the set of pairs to a size such that more computationally expensive comparisons can be made. In §4.1, we apply these concepts.

### 2.2 Cluster-Based Blocking

Others have described blocking as a clustering problem, sometimes with a special emphasis on privacy, e.g., see [7, 13, 14, 21]. The motivation is natural: the records in a cluster should be similar, making good candidate pairs for linkage.

One clustering approach proposed for blocking is nearest neighbor clustering. Threshold nearest neighbor clustering (TNN) begins with a single record as the base of the first cluster, and recursively adds the nearest neighbors of records in the cluster until the distance<sup>1</sup> to the nearest neighbor exceeds some threshold. Then one of the remaining records is picked to be the base for the next cluster, and so forth.  $k$ -nearest neighbor clustering (KNN) uses a similar procedure, but ensures that each cluster contains at least  $k$  records<sup>2</sup>, to help maintain “ $k$ -anonymity” [13]. A major drawback of nearest neighbor clustering is that it requires computing a large number of distances between records,  $O(n^2)$ . Blocking a new record means finding its nearest neighbors, an  $O(n)$  operation.

---

<sup>1</sup> The distance metric used can vary depending on the nature of the records.

<sup>2</sup> Privacy-preserving versions of these approaches use “reference values” rather than the records themselves to cluster the records [21].

The cost of calculating distances between records in large, high-dimensional datasets led [16] to propose the method of *canopies*. In this approach, a computationally cheap (if inaccurate) distance metric is used to place records into potentially-overlapping sets (canopies). An initial record is picked randomly to be the base of the first canopy; all records within a distance  $t_1$  of the base are grouped under that canopy. Those within distance  $t_2 \leq t_1$  of the base are removed from later consideration. A new record is picked to be the base of the next canopy, and the procedure is repeated until the list of candidate records is empty. More accurate but expensive distance measures are computed only between records that fall under at least one shared canopy. That is, only record-pairs sharing a canopy are candidates to be linked.

Canopies is not strictly a blocking method. They overlap, making the collection of canopies only a covering of the set of records, rather than a partition. We can derive blocks from canopies, either set-theoretically or by setting  $t_1 = t_2$ . The complexity of building the canopies is  $O(nC_n)$ , with  $C_n$  being the number of canopies, itself a complicated and random function of the data, the thresholds, and the order in which records are chosen as bases. Further, finding fast, rough distance measures for complicated high-dimensional records is non-trivial.

## 2.3 LSH-Based Approaches

We explore two LSH-based blocking methods. These are based, respectively, on graph partitioning or community discovery, and on combining random projections with classical clustering. The main reason for exploring these two methods is that even with comparatively efficient algorithms for partitioning the similarity graph, doing that is still computationally impractical for hundreds of thousands of records.

### 2.3.1 Shingling

LSH-based blocking schemes “shingle” [19] records. That is, each record is treated as a string and is replaced by a “bag” (or “multi-set”) of length- $k$  contiguous sub-strings that it contains. These are known as “ $k$ -grams”, “shingles”, or “tokens”. For example, the string “TORONTO” yields the bag of length-two shingles “TO”, “OR”, “RO”, “ON”, “NT”, “TO”. (N.B., “TO” appears twice.)

As alternative to shingling, we might use a bag-of-words representation, or even to shingle into consecutive pairs (triples, etc.) of words. In our experiments, shingling at the level of letters worked better than dividing by words.

### 2.3.2 Transitive LSH (TLSH)

We create a graph of the similarity between records. For simplicity, assume that all fields are string-valued. Each record is shingled with a common  $k$ , and the bags of shingles for all  $n$  records are reduced to an  $n$ -column binary-valued matrix  $M$ , indicating which shingles occur in which records.  $M$  is large, since the number of length- $k$  shingles typically grows exponentially with  $k$ . As most shingles are absent from most records,  $M$  is sparse. We reduce its dimension by

generating a random “minhash” function and applying it to each column. Such functions map columns of  $M$  to integers, ensuring that the probability of two columns being mapped to the same value equals the Jaccard similarity between the columns [19]. Generating  $p$  different minhash functions, we reduce the large, sparse matrix  $M$  to a dense  $p \times n$  matrix,  $M'$ , of integer-valued “signatures,” while preserving information. Each row of  $M'$  is a random projection of  $M$ . Finally, we divide the rows of  $M'$  into  $b$  non-overlapping “bands,” apply a hash function to each band and column, and establish an edge between two records if their columns of  $M'$  are mapped to the same value in any band.<sup>3</sup>

These edges define a graph: records are nodes, and edges indicate a certain degree of similarity between them. We form blocks by dividing the graph into its connected components. However, the largest connected components are typically very large, making them unsuitable as blocks. Thus, we sub-divide the connected components into “communities” or “modules” — sub-graphs that are densely connected internally, but sparsely connected to the rest of the graph. This ensures that the blocks produced consist of records that are all highly similar, while having relatively few ties of similarity to records in other blocks [8]. Specifically, we apply the algorithm of [6]<sup>4</sup>, sub-dividing communities greedily, until even the largest community is smaller than a specified threshold.<sup>5</sup> The end result is a set of blocks that balance false negative errors in linkage (minimized by having a few large blocks) and the speed of linkage (minimized by keeping each block small). We summarize the whole procedure in Algorithm 2.3.2 (see Appendix 5).

TLSH involves many tuning parameters (the length of shingles, the number of random permutations, the maximum size of communities, etc.) We chose the shingle such that we have the highest recall possible for each application. We used a random permutation of 100, since the recall was approximately constant for all permutations higher than 100. Furthermore, we chose a maximum size of the communities of 500, after tuning this specifically for desired speed.

### 2.3.3 K-Means Locality Sensitive Hashing (KLSH)

The second LSH-based blocking method begins, like TLSH, by shingling the records, treated as strings, but then differs in several ways. First, we do not ignore the number of times each shingle type appears in a record, but rather keep track of these counts, leading to a bag-of-shingles representation for records. Second, we measure similarity between records using the inner product of bag-of-shingles vectors, with inverse-document-frequency (IDF) weighting. Third, we reduce the dimensionality of the bag-of-shingles vectors by random projections, followed by clustering the low-dimensional projected vectors with the  $k$ -means algorithm. Hence, we can control the mean number of records per cluster to be  $n/c$ , where  $c$  is the number of block-clusters. In practice, there is a fairly small dispersion around this mean, leading to blocks that, by construction, have the roughly

---

<sup>3</sup> To be mapped to the same value in a particular band, two columns must either be equal, or a low-probability “collision” occurred for the hash function.

<sup>4</sup> We could use other community-discovery algorithms, e.g. [9].

<sup>5</sup> This maximum size ensures that record linkage is feasible.

the same distribution for all applications.<sup>6</sup> The KLSH algorithm is given in Appendix 5.

### 3 Computational Complexity

#### 3.1 Computational Complexity of TLSH

The first steps of the algorithm can be done independently across records. Shingling a single record is  $O(1)$ , so shingling all the records is  $O(n)$ . Similarly, applying one minhash function to the shingles of one record is  $O(1)$ , and there are  $p$  minhash functions, so minhashing takes  $O(np)$  time. Hashing again, with  $b$  bands, takes  $O(nb)$  time. We assume that  $p$  and  $b$  are both  $O(1)$  as  $n$  grows.

We create an edge between every pair of records that get mapped to the same value by the hash function in some band. Rather than iterating over pairs of records, it is faster to iterate over values  $v$  in the range of the hash function. If there are  $|v|$  records mapped to the value  $v$ , creating their edges takes  $O(|v|^2)$  time. On average,  $|v| = nV^{-1}$ , where  $V$  is the number of points in the range of the hash function, so creating the edge list takes  $O(V(n/V)^2) = O(n^2V^{-1})$  time. [6] shows that creating the communities from the graph is  $O(n(\log n)^2)$ .

The total complexity of TLSH is  $O(n) + O(np) + O(nb) + O(n^2V^{-1}) + O(n(\log n)^2) = O(n^2V^{-1})$ , and is dominated by actually building the graph.

#### 3.2 Computational Complexity of KLSH

As with TLSH, the shingling phase of KLSH takes  $O(n)$  time. The time required for the random projections, however, is more complicated. Let  $w(n)$  be the number of distinct words found across the  $n$  records. The time needed to do one random projection of one record is then  $O(w(n))$ , and the time for the whole random projection phase is  $O(npw(n))$ . For  $k$ -means cluster, with a constant number of iterations  $I$ , the time required to form  $b$  clusters of  $n$   $p$ -dimensional vectors is  $O(bnpI)$ . Hence, the complexity is  $O(npw(n)) + O(bnpI)$ .

Heaps's law suggests  $w(n) = O(n^\beta)$ , where  $0 < \beta < 1$ .<sup>7</sup> Thus, the complexity is  $O(pn^{1+\beta}) + O(bnpI)$ . For record linkage to run in linear time, it must run in constant time in each block. Thus, the number of records per block must be constant, i.e.,  $b = O(n)$ . Hence, the time-complexity for blocking is  $O(pn^{1+\beta}) + O(n^2pI) = O(n^2pI)$ , a quadratic time algorithm dominated by the clustering. Letting  $b = O(1)$  yields an over-all time complexity of  $O(pn^{1+\beta})$ , dominated by the projection step. If we assume  $\beta = 0.5$  and let  $b = O(\sqrt{n})$ , then both the projection and the clustering steps are  $O(pn^{1.5})$ . Record linkage in each block is  $O(n)$ , so record linkage is  $O(n^{1.5})$ , rather than  $O(n^2)$  without blocking.

<sup>6</sup> This property is not guaranteed for most LSH methods.

<sup>7</sup> For English text,  $0.4 < \beta < 0.6$ .

### 3.3 Computational Complexity of Traditional Blocking Approaches

Traditional blocking approaches use attributes of the records to partition records into blocks. As such, calculating the blocks using traditional approaches requires  $O(n)$  computations. For example, approaches that block on birth year only require a partition of the records based on these fields. That is, each record is simply mapped to one of the unique birth year values in the dataset, which is an  $O(n)$  calculation for a list of size  $n$ . Some traditional approaches, however, require  $O(n^2)$  computations. For example, in Table 1, we show some effective blocking strategies which require  $O(n^2)$  computations, but each operation is so cheap that they can be run in reasonable time for moderately sized files.

## 4 Results

We test the previously mentioned approaches on data from the RecordLinkage R package.<sup>8</sup> These simulated datasets contain 500 and 10,000 records (denoted `RLdata500` and `RLdata10000`), with exactly 10% duplicates in each list. These datasets contain first and last Germanic name and full date of birth (DOB). Each duplicate contains one error with respect to the original record, and there is maximum of one duplicate per original record. Each record has a unique identifier, allowing us to test the performance of the blocking methods.

We explore the performance of the previously presented methods under other scenarios of measurement error. [1, 4, 5] developed a data generation and corruption tool that creates synthetic datasets containing various field attributes. This tool includes dependencies between fields and permits the generation of different types of errors. We now describe the characteristics of the datafiles used in the simulation. We consider three files having the following field attributes: first and last name, gender, postal code, city, telephone number, credit card number, and age. For each database, we allow either 10, 30, or 50% duplicates per file, and each duplicate has five errors with respect to the original record, where these five errors are allocated at random among the fields. Each original record has maximum of five duplicates. We refer to these files as the “noisy” files.

### 4.1 Traditional Blocking Approaches

Tables 1 – 2 provide results of traditional blocking when applied to the `RLdata10000` and “noisy” files. While field-specific information *can* yield favorable blocking solutions, each blocking criteria is application specific. The overall goal of blocking is to reduce the overall set of candidate pairs, while minimizing the false negatives induced. Thus, we find the *recall* and *reduction ratio* (RR). This corresponds to the proportion of true matches that the blocking criteria preserves, and the proportion of record-pairs discarded by the blocking, respectively.

---

<sup>8</sup> <http://www.inside-r.org/packages/cran/RecordLinkage/docs/RLdata>

Criteria 1 – 5 (Table 1) and 1 – 6 (Table 2) show that *some* blocking approaches are poor, where the recall is never above 90%. Criteria requiring exact agreement in a single field or on a combination of them are susceptible to field errors. More reliable criteria are constructed using combinations of fields such that multiple disagreements must be met for a pair to be declared as a non-match. (See Criteria 7–10 and 12 in Table 1, and 7 – 8 in Table 2.) We obtain high recall and RR using these, but in general their performance is context-dependent.

Criteria 10 (Table 1) deals with the case when a pair is declared a non-match whenever it disagrees in four or more fields, which is reliable since false-negative pairs are only induced when the datafile contains large amounts of error. For example this criterion does not lead to good results with the noisy files, hence a stronger criteria is needed, such as 7 (Table 2). Using Criteria 12 (Table 1) and 8 (Table 2), we further reduce the set of candidate pairs whenever a pair has a strong disagreement in an important field.<sup>9</sup> These criteria are robust. In order to induce false negatives, the error in the file must be much higher than expected.

**Table 1.** Criteria for declaring pairs as non-matches, where results correspond to the RLdata10000 datafile

Declare non-match if disagreement in:	Recall (%)	RR (%)
1. First OR last name	39.20	99.98
2. Day OR month OR year of birth	59.30	99.99
3. Year of birth	84.20	98.75
4. Day of birth	86.10	96.74
5. Month of birth	88.40	91.70
6. Decade of birth	93.20	87.76
7. First AND last name	99.20	97.36
8. {First AND last name} OR {day AND month AND year of birth}	99.20	99.67
9. Day AND month AND year of birth	100.00	87.61
10. More than three fields	100.00	99.26
11. Initial of first OR last name	100.00	99.25
12. {More than three fields} OR {Levenshtein dist. $\geq 4$ in first OR last name}	100.00	99.97

**Table 2.** Criteria for declaring pairs as non-matches, where results correspond to the noisy datafile with 10% duplicates. Similar results obtained for 30 and 50% duplicates.

Declare non-match if disagree in:	Recall (%)	RR (%)
1. Gender	31.96	53.39
2. City	31.53	77.25
3. Postal Code	32.65	94.20
4. First OR last name	1.30	>99.99
5. Initial of first OR last name	78.10	99.52
6. First AND last name	28.97	99.02
7. All fields	93.28	40.63
8. {All fields} OR {Levenshtein dist. $\geq 4$ in first OR last name}	92.84	99.92

## 4.2 Clustering Approaches


Our implementations of [16]’s canopies approach and [21]’s nearest neighbor approach perform poorly on the RLdata10000 and “noisy” datasets<sup>10</sup>. Figure 1 gives results of these approaches for different threshold parameters ( $t$  is the threshold parameter for sorted TNN) for the RLdata10000 dataset. For all thresholds, both TNN and canopies fail to achieve a balance of high recall and a high reduction ratio.

Turning to the “noisy” dataset with 10% duplicates, we find that TNN fails to achieve a balance of high recall and high reduction ratio, regardless of the


<sup>9</sup> We use the Levenshtein distance (LD) of first and last names for pairs passing Criterion 10 of Table 1 or Criteria 7 of Table 2, and declare pairs as non-matches when  $LD \geq 4$  in either first or last name.

<sup>10</sup> In our implementations, we use the TF-IDF matrix representation of the records and Euclidean distance to compare pairs of records in TNN and canopies. We tried several other distance measures, each of which gave similar results.





**Fig. 1.** Performance of threshold nearest neighbors (left) and canopies (right) on the RLdata10000 datafile



**Fig. 2.** Performance of TNN (left) and canopies (right) on the “noisy” datafile (10% duplicates). The other “noisy” datafiles exhibited similar behavior as the figures above.

threshold  $t$  that is used. Similarly, the canopies approach does not yield a balance of high recall while reducing the number of candidate pairs.

Clearly, both clustering approaches fail to achieve a balance of high recall and RR for any threshold parameters. The inefficacy of these approaches is likely due the limited number of field attributes (five fields) and the Euclidean distance metric used for these datasets. In particular, only three fields in the “noisy” dataset use textual information, which both of these approaches use to identify similar records. Limited field information can make it difficult for clustering approaches to group similar records together, since the resulting term frequency matrices will be very sparse. Thus, we investigate the behavior with the same number of duplicates, but vary the error rate and provide richer information at the field attribute level. Figure 2 illustrates that both methods do not have a good balance between recall and RR, which we investigated for various thresholds. As

such, further analysis of these approaches on more information-rich datasets is required in order to make sound conclusions about their efficacy for blocking. (We note that the metrics used in TLSH and KLSH, which shingle the records, were chosen so as to not have such problems.)

### 4.3 LSH Approaches

Since the performance of KLSH and TLSH depends on tuning parameters, we tune each application appropriately to these. We empirically measure the scalability of these methods, which are consistent with our derivations in §3.

We analyze the `RLdata10000` database for TLSH and KLSH. As we increase  $k$  under TLSH, we see that the recall peaks at  $k = 5$ , and does very poorly (below 40% recall) when  $k \leq 4$ . For KLSH, the highest and most consistent recall is when  $k = 2$ , since it is always above 80% and it is about the same no matter the total number of blocks chosen (see Figure 4). In terms of RR, we see that TLSH performs extremely poorly as the total number of blocks increases, whereas KLSH performs extremely well in terms of RR comparatively (Figure 5). Figure 3 shows empirically that the running time for both KLSH and TLSH scales quadratically with the  $n$ , matching our asymptotic derivation. We then analyze the “noisy” database for TLSH and KLSH (see Figures 6 and 7).

#### 4.3.1 Comparisons of Methods

In terms of comparing to the methods presented in Table 1, we find that TLSH is not comparable in terms of recall or RR. However, KLSH easily beats Criteria 1–2 and competes with Criteria 3–4 on both recall and RR. It does not perform as well in terms of recall as the rest of the criteria, however, it *may* in other applications with more complex information for each record (this is a subject of future work). When comparing the Table 2 to TLSH and KLSH when run for the noisy datafile, we find that TLSH and KLSH usually do better when tuned properly, however not always. Due to the way these files have been constructed, more investigation need to be done in terms of how naive methods work for real work type applications versus LSH-based methods.

Comparing to other blocking methods, both KLSH and TLSH outperform KNN in terms of recall (and RR for the noisy datafiles). We find that for this dataset, canopies do not perform well in terms of recall or RR unless a specific threshold  $t_1$  is chosen. However, given this choice of  $t_1$ , this approach yields either high recall and low RR or vice versa, making canopies undesirable according to our criteria.

For the `RLdata10000` dataset, the simple yet effective traditional blocking methods and KLSH perform best in terms of balancing both high recall and high RR. As already stated, we expect the performance of these to be heavily application-dependent. Additionally, note that each method relies on high-quality labeled record linkage data to measure the recall and RR and the clustering methods require tuning parameters, which can be quite sensitive. Our studies show that TLSH is the least sensitive in general and further explorations should

be done here. Future work should explore the characteristics of the underlying datasets for which one method would be preferred over another.

### 4.3.2 Sensitivity Analysis on RLdata500 and RLdata10000


A sensitivity analysis is given for KLSH and TLSH. For TLSH, the RLdata500 dataset is not very sensitive to  $b$  since the recall is always above 80% whereas the RLdata10000 dataset is quite sensitive to the band, and we recommend the use of a band of 21–22 since the recall for these  $b$  is  $\approx 96\%$ , although this may change for other datasets. We then evaluate TLSH using the “best” choice of the band for shingled values from  $k = 1, \dots, 5$ . The sensitivity analysis for the “noisy” datafiles was quite similar to that described above, where a band of 22 was deemed the most appropriate for TLSH. For KLSH, we found that we needed to increase the number of permutations slightly to improve the recall and recommend  $p = 150$ .

For KLSH, we find that when the number of random permutations  $p$  is above 100, the recall does not change considerably. We refer back to Figure 4 (right), which illustrates the recall versus number of blocks when  $p = 100$ . When  $k = 4$ , the recall is always above 70%. However, we find that when  $k = 2$ , the recall is always above 80%.


## 5 Discussion

We have explored two LSH methods for blocking, one of which would naturally fit into the privacy preserving record linkage (PPRL) framework, since the method could be made to be private by creating reference values for each individual in the database. This has been done for many blocking methods in the context of PPRL [7, 13, 14, 22]. KLSH performs just as well or better than commonly used blocking methods, such as some simple traditional blocking methods, nearest neighbor clustering approaches, and canopies [16, 21]. One drawback is that like LSH-based methods, it must be tuned for each application since it is sensitive to the tuning parameters. Thus, some *reliable* training data must be available to evaluate the recall and RR (and tune KLSH or clustering type methods). In many situations, a researcher may be better off by using domain-specific knowledge to reduce the set of comparisons, as shown in §4.1.

LSH-methods have been described elsewhere as “private blocking” due to the hashing step. However, they do not in fact provide any formal privacy guarantees in our setting. The new variant that we have introduced, KLSH, does satisfy the  $k$ -anonymity criterion for the de-duplication of a single file. However, the data remain subject to intruder attacks, as the literature on differential privacy makes clear, and the vulnerability is greater the smaller the value of  $k$ . Our broader goal, however, is to merge and analyze data from multiple files. Privacy protection in that context is far more complicated. Even if one could provide privacy guarantees for each file separately, it would still be possible to identify specific entities or *sensitive* information regarding entities in the merged database.




**Fig. 3.** RLdata10000/RLdata500 datasets. Left: Square Root Elapsed time versus number of records for KLSH and TLSH, illustrating that both methods scale nearly quadratically (matching the computationally complexity findings). We shingle using  $k = 5$  for both methods. We use a band of 26 for TLSH. Right: Recall versus  $b$  for both RLdata500 and RLdata10000 after running TLSH.




**Fig. 4.** RLdata10000 dataset. Left: Recall versus number of shingles  $k$  for KLSH. The highest recall occurs at  $k = 5$ . Right: Recall versus the total number of blocks, where we vary the number of shingles  $k$ . We find that the highest recall is for  $k = 2$ .


The approach of PPRL reviewed in [10] sets out to deal with this problem. Merging data from multiple files with the same or similar values without releasing their attributes is what PPRL hopes to achieve. Indeed, one of course needs to go further, since performing statistical analyses on the merged database is the real objective of PPRL. Whether the new “private blocking” approaches discussed offer any progress on this problem, it is unclear at best. Adequately addressing the PPRL goals remains elusive, as do formal privacy guarantees, be they from differential privacy or other methods.



**Fig. 5.** RLdata10000 dataset. Left: For TLSH, we see the RR versus the number of shingles, where the RR is always very high. We emphasize that TLSH does about as well on the RR as any of the other methods, and certainly does much better than many traditional blocking methods and KNN. (The RR is always above 98% for all shingles with  $b = 26$ .) Right: For KLSH, we illustrate the RR versus the total number of blocks for various  $k = 1, \dots, 4$  illustrating that as the number of blocks increases, the RR increases dramatically. When the total block size is at least 25, the  $RR \geq 95\%$ .



**Fig. 6.** Left: We run TLSH for 10 percent duplicates, as before, the application is quite sensitive to  $b, k$ . Hence, it is quite easy to find values of  $b, k$  such that the recall is very low or if tuning is done properly, we can find values of  $b, k$  where the recall is acceptable. We note this relies on very good ground truth. The only value of  $k$  we recommend is 4 since it is close to 90% recall. The computational time is the same as previously. Right: Elapsed time for 10, 30, and 50 percent duplicates on “noisy” dataset.



**Fig. 7.** We run KLSH at 10 percent duplicates with  $p=100$  (left) and  $p=150$  (right). We see as the number of permutations increases (left figure), the recall increases. The behavior is the same for 30 and 50 percent duplicates. This indicates that KLSH needs to be tuned for each application based on  $p$ . (We note that KLSH was run here using a bag of words type formulation. Shingling should be explored in future work.)

**Acknowledgements.** We thank Peter Christen, Patrick Ball, and Cosma Shalizi for thoughtful conversations that led to early versions of this manuscript. We also thank the reviewers for their suggestions and comments.

## References

- [1] Christen, P.: Probabilistic Data Generation for Deduplication and Data Linkage. In: Gallagher, M., Hogan, J.P., Maire, F. (eds.) IDEAL 2005. LNCS, vol. 3578, pp. 109–116. Springer, Heidelberg (2005)
- [2] Christen, P.: A survey of indexing techniques for scalable record linkage and deduplication. *IEEE Transactions on Knowledge and Data Engineering* 24 (2012)
- [3] Christen, P., Pudjijono, A.: Accurate synthetic generation of realistic personal information. In: Theeramunkong, T., Kijssirikul, B., Cercone, N., Ho, T.-B. (eds.) PAKDD 2009. LNCS, vol. 5476, pp. 507–514. Springer, Heidelberg (2009)
- [4] Christen, P., Pudjijono, A.: Accurate Synthetic Generation of Realistic Personal Information. In: Theeramunkong, T., Kijssirikul, B., Cercone, N., Ho, T.-B. (eds.) PAKDD 2009. LNCS, vol. 5476, pp. 507–514. Springer, Heidelberg (2009), [http://dx.doi.org/10.1007/978-3-642-01307-2\\_47](http://dx.doi.org/10.1007/978-3-642-01307-2_47)
- [5] Christen, P., Vatsalan, D.: Flexible and Extensible Generation and Corruption of Personal Data. In: Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM 2013) (2013)
- [6] Clauset, A., Newman, M.E., Moore, C.: Finding community structure in very large networks. *Physical Review E* 70, 066111 (2004)
- [7] Durham, E.A.: A framework for accurate, efficient private record linkage. Ph.D. thesis, Vanderbilt University (2012)
- [8] Fortunato, S.: Community detection in graphs. *Physics Reports* 486, 75–174 (2010)

- [9] Goldenberg, A., Zheng, A.X., Fienberg, S.E., Airoidi, E.M.: A survey of statistical network models. *Foundations and Trends® in Machine Learning* 2, 129–233 (2010)
- [10] Hall, R., Fienberg, S.: Valid statistical inference on automatically matched files. In: Domingo-Ferrer, J., Tinnirello, I. (eds.) *PSD 2012. LNCS*, vol. 7556, pp. 131–142. Springer, Heidelberg (2012)
- [11] Herzog, T., Scheuren, F., Winkler, W.: *Data Quality and Record Linkage Techniques*. Springer, New York (2007)
- [12] Herzog, T., Scheuren, F., Winkler, W.: Record linkage. *Wiley Interdisciplinary Reviews: Computational Statistics* 2 (2010), doi:10.1002/wics.108
- [13] Karakasidis, A., Verykios, V.S.: Reference table based k-anonymous private blocking. In: *Proceedings of the 27th Annual ACM Symposium on Applied Computing*, pp. 859–864. ACM (2012)
- [14] Kuzu, M., Kantarcioglu, M., Durham, E., Malin, B.: A constraint satisfaction cryptanalysis of bloom filters in private record linkage. In: Fischer-Hübner, S., Hopper, N. (eds.) *PETS 2011. LNCS*, vol. 6794, pp. 226–245. Springer, Heidelberg (2011)
- [15] Liang, H., Wang, Y., Christen, P., Gayler, R.: Noise-tolerant approximate blocking for dynamic real-time entity resolution. In: Tseng, V.S., Ho, T.B., Zhou, Z.-H., Chen, A.L.P., Kao, H.-Y. (eds.) *PAKDD 2014, Part II. LNCS*, vol. 8444, pp. 449–460. Springer, Heidelberg (2014)
- [16] McCallum, A., Nigam, K., Ungar, L.H.: Efficient clustering of high-dimensional data sets with application to reference matching. In: *Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 169–178. ACM (2000)
- [17] Pasula, H., Marthi, B., Milch, B., Russell, S., Shpitser, I.: Identity uncertainty and citation matching. In: *Advances in Neural Information Processing Systems*, pp. 1425–1432 (2003)
- [18] Paulevé, L., Jégou, H., Amsaleg, L.: Locality sensitive hashing: A comparison of hash function types and querying mechanisms. *Pattern Recognition Letters* 31, 1348–1358 (2010)
- [19] Rajaraman, A., Ullman, J.D.: *Mining of massive datasets*. Cambridge University Press (2012)
- [20] Steorts, R.C., Hall, R., Fienberg, S.: A Bayesian approach to graphical record linkage and de-duplication (2013) (submitted)
- [21] Vatsalan, D., Christen, P.: Sorted nearest neighborhood clustering for efficient private blocking. In: Pei, J., Tseng, V.S., Cao, L., Motoda, H., Xu, G. (eds.) *PAKDD 2013, Part II. LNCS*, vol. 7819, pp. 341–352. Springer, Heidelberg (2013)
- [22] Vatsalan, D., Christen, P., Verykios, V.S.: An efficient two-party protocol for approximate matching in private record linkage. In: *Proceedings of the Ninth Australasian Data Mining Conference*, vol. 121, pp. 125–136. Australian Computer Society, Inc. (2011)
- [23] Winkler, W., Yancey, W., Porter, E.: Fast record linkage of very large files in support of decennial and administrative records projects. In: *Proceedings of American Statistical Association Section on Survey Research Methods* (2010)

## Algorithms for KLSH and TLSH

We provide the algorithms for KLSH and TLSH below (see §2.3.2 and 2.3.3):

---

### Algorithm 1. Transitive Locality Sensitive Hashing (TLSH)

---

Place similar records into blocks and impose transitivity

**Data:**  $X_{ij}$ , tuning parameters  $b, t, k$

Shingle each  $X_{ij}$  into length- $k$  strings

Create a binary matrix  $M$  indicating which tokens appear in which records

Create an integer-valued matrix  $M'$  of minhash signatures from  $M$

Divide the rows of  $M'$  into  $b$  bands

**for each band do**

- | Apply a random hash function to the band of  $M'$
- | Record an edge between two records if the hash maps them to the same bucket

**end**

**while the largest community has  $> t$  records do**

- | Cut the edge graph into finer communities using the algorithm of [6]

**end**

**return the final list of communities**

---



---

### Algorithm 2. K-Means Locality Sensitive Hashing (KLSH). The number of blocks $c$ is set by $c = n / (\text{desired avg. number of records per block})$ .

---

Place similar records into blocks and using k-means clustering and random projections

**Data:**  $X_{ij}$ , number of desired blocks  $c$ , tokenization tuning parameters  $\tau$ , number of projections  $p$

**for each record  $X_{ij}$  do**

- | **Set**  $v_{ij} = \text{TOKENIZE}(X_{ij}, \tau)$

**end**

**for each token  $w$  do**

- | **Set**  $N_w = \text{number of bags containing } w$
- | **Set**  $IDF_w = \log n / N_w$

**end**

**for  $m$  from 1 to  $p$  do**

- | **Set**  $u_m = \text{a random unit vector}$
- | **for each bag-of-tokens vector  $v_{ij}$  do**

  - | | **Set**  $r_{ijm} = \sum_w u_{iw} v_{ijw} IDF_w$

- | **end**

**end**

**return**  $\text{KMEANS}(r, c)$

---