Em-K Indexing for Approximate Query Matching in Large-scale ER

SAMUDRA HERATH, University of Adelaide, Australia

MATTHEW ROUGHAN, University of Adelaide, Australia

GARY GLONEK, University of Adelaide, Australia

Accurate and efficient entity resolution (ER) is a significant challenge in many data mining and analysis projects requiring integrating and processing massive data collections. It is becoming increasingly important in real-world applications to develop ER solutions that produce prompt responses for entity queries on large-scale databases. Some of these applications demand entity query matching against large-scale reference databases within a short time. We define this as the query matching problem in ER in this work. Indexing or blocking techniques reduce the search space and execution time in the ER process. However, approximate indexing techniques that scale to very large-scale datasets remain open to research. In this paper, we investigate the query matching problem in ER to propose an indexing method suitable for approximate and efficient query matching.

We first use spatial mappings to embed records in a multidimensional Euclidean space that preserves the domain-specific similarity. Among the various mapping techniques, we choose multidimensional scaling. Then using a Kd-tree and the nearest neighbour search, the method returns a block of records that includes potential matches for a query. Our method can process queries against a large-scale dataset using only a fraction of the data L (given the dataset size is N), with a $O(L^2)$ complexity where $L \ll N$. The experiments conducted on several datasets showed the effectiveness of the proposed method.

CCS Concepts: • **Information systems** → **Data cleaning**; **Entity resolution**; **Deduplication**.

Additional Key Words and Phrases: Data integration, data linkage, data matching, large-scale data, record linkage.

1 INTRODUCTION

Many real-world data collections are low in quality because of errors (e.g., typographical errors or phonetic misspellings), incomplete or missing data, incompatible formats for recording database fields (e.g., dates, addresses), and temporal inconsistencies. Integrating or querying multiple data sources to identify records that belong to the same real-world entity is a challenging task in the presence of such data. This task is often referred to as the *entity resolution (ER)* problem and appears in many database applications when identifying duplicates, cleansing data, or improving data quality.

ER has been studied for years. However, the issues of efficiently handling large-scale data remain an open research problem [22]. Typical ER algorithms have a quadratic running time, which is computationally prohibitive for large-scale data collections. This performance bottleneck occurs due to the detail-level pairwise comparison step of the ER process.

Consider the motivating example of linking medical records where one entity can have many entries over several years. A public health case study could involve hundreds of thousands of entities with millions of records. Quadratic ER among the records will require trillions of pairwise comparisons in this application, which is computationally infeasible.

Indexing techniques address this problem by grouping similar records into blocks. Usually, groups of similar records within or between datasets are smaller than the total number of records.

Hence, many comparisons will be among the non-matching records when applying pairwise comparisons using a brute-force ER algorithm. Indexing techniques aim to reduce the potential number of comparisons by reducing the comparisons between those non-matches [8]. Therefore, indexing-based ER solutions run much faster than brute-force ER solutions [18].

However, most existing indexing methods still have quadratic time complexity and are too slow to deal with very large-scale data collections. Also, some non-scaling ER algorithms tend to generate a large number of blocks and a large number of candidate records in each block for large-scale scenarios [8].

Traditional ER solutions often process databases offline in batch mode, and no further action is required once a pair of matches are determined. However, many organisations are moving online where they have to provide their services through prompt responses. Hence, many newer, real-world scenarios require real-time query processing against large-scale databases. Some of the applications demand entity query matching against large-scale reference databases within a short time. We refer to this as the *query matching* problem in ER. In such configurations, traditional ER solutions become less efficient or unusable.

In order to motivate the problem context and illustrate the usefulness of the approach presented in this paper, we provide the following real-world example. The application is in criminal investigations, where law enforcement officers need to query a database to identify potential suspects. Suspects usually lie to police investigators about their true identities, e.g., names, birth dates, or addresses [6]. Hence, finding an exact match for falsified data against the real identity recorded in a law enforcement computer system is problematic. Detecting deceptive identities is a time-consuming activity that involves large amounts of manual information processing in real-life scenarios [32]. Therefore efficient ER solutions that support real-time and approximate query matching against existing datasets can be invaluable.

Toward the challenge of real-time approximate query matching, we present an indexing technique that reduces the number of pairwise comparisons needed in the ER. The proposed indexing method transforms a set of records into a set of vectors in a metric-space, specifically a lower-dimensional Euclidean space using multidimensional scaling. These vectors have two main attractive attributes in the context. First, comparisons between vectors in a metric-space are much cheaper than string comparisons. Second, these vectors support efficient indexing data structures. We utilise these properties to propose an indexing approach that classifies similar vectors in a low-dimensional Euclidean space. We call our method *Em-K indexing*, as our method operates by embedding data in a K-dimensional space.

Exact query matching is easy against a reference database as we can search based on lexicographical order using an efficient data structure, e.g., binary trees. However, exact query matching is impossible with many real-world databases due to various data quality issues, and approximate query matching is the common approach [7]. The proposed indexing method address the problem of real-time and approximate query matching. We will be mainly using the term query matching instead of approximate query matching in this paper.

Our main objective is to develop a fast indexing method for query matching. We want our method to be efficient for large-scale data and robust to errors in the data. The proposed method searches for a block of records in the reference database as a set of potential matches to the querying record. Hence we avoid comparing the query against all the records in the referencing database. However, the search is done in the Euclidean space using a Kd-tree data structure where each record in the reference database and a query record become a vector in this space.

Given the block size is B, we search for this number of nearest neighbours for the query point by traversing the Kd-tree nodes that store the reference database. The search has a $O(B \log N)$ complexity, given N is the number of records in the reference database. The method embeds query

records in a pre-mapped Euclidean space to search against the reference data points. We propose an out-of-sample embedding that uses a fraction of the original data (defined as landmarks in Section 5) for query record embedding. For L landmarks, the embedding requires $O(L^2)$ operations, and we can choose L such that $L \ll N$.

In this paper, we explore the use of metric-space indexing for efficient and approximate query matching. In particular, our contributions are,

- We formulate the query matching and deduplication problems in ER to provide metric-space base solutions. First, we propose an indexing approach based on landmarks as a motivating example to explore the basic building blocks needed for query matching.
- We propose a landmarks based indexing technique for query matching to provide a quick and accurate block of potential matches. Our method can process a stream of queries against a large-scale data set within a short time. By doing so, we obtain as many of the matching records as possible where the processing time of a single query takes a sub-second time. The technique is robust to noisy data that contains errors and allow efficient approximate query matching.

2 PRELIMINARIES

To describe the problem succinctly, we first describe relevant definitions and some key concepts. The definitions of relevant concepts here follow those in [26, 29].

2.1 Entity Resolution

An *entity* is a real-world object, e.g., person, place or product, that has a unique identifier to distinguish it from other entities of the same type. An *entity profile* describes an entity using a collection of name-value pairs. A set of entity profiles is called *entity collection*, denoted by E. A pair of similar entity profiles are called *duplicates*. A duplicate of an entity can be either an exact copy of the original entity profile or an entity profile that contains an error (e.g., typographical error). A database representation of an entity profile is usually referred to as a *record*.

DEFINITION 1. *(Entity Resolution): Given two records r_i, r_j is a match, if they refer to the same unique real-world object. We denote this as $r_i \equiv r_j$. The goal of ER is to link different records that describe the same entity within an entity collection or across two or more entity collections.*

2.2 ER Tasks

Following the above definitions, we distinguish between the following ER tasks.

(1) *Dirty ER:* Given an entity collection E that contains entities $e_1, e_2, ..., e_n$, find all duplicates and produces a set of equivalence clusters of distinct entities. It is also known as deduplication in many database applications.

(2) *Clean−Clean ER:* Given two duplicate free entity collection E_1 and E_2 that contains entities $e_1, e_2, ..., e_n$, find all records that belong to a single entity. Our indexing method is aimed at the *Clean−Clean ER* problem where one dataset is a stream of queries, and the other is a reference database. However, we use deduplication as a motivating example that explains the relevant building blocks of the proposed method.

2.3 Record Comparisons

Similar records pairs of an entity are determined by applying a similarity function over the corresponding attributes of two records. Assume a pair of records (r_i, r_j) and a set of attributes $a_1, a_2, ..., a_x$ that describe them. The similarities $s_1, ..., s_x$ between attribute values are determined by applying a set of similarity functions $sim(r_i.a_k, r_j.a_k)$, with $1 \leq k \leq x$ for each pair of attribute

values. Then a total similarity score of $S = \sum_{i=1}^{x} s_i$ is calculated to classify the record pairs as a match or a non-match based on a matching threshold.

Several comparison methods such as edit distance, a.k.a Levenshtein distance, Jaro distance, and q-gram distance are found in the domain of strings [23]. In this work, we mainly used Levenshtein distance to measure the similarity at the attribute level. It calculates the minimum number of character insertions, deletions, and replacements necessary to transform a string s_1 into a string s_2. Minkowski metrics based on L^p norms, $\| x \|_p = (\sum_i |x_i|^p)^{1/p}$, with $p \geqslant 1$ are a common used class of vector spaces. For our vector space, we used the most common Minkowski metric; Euclidean distances $d_E(p = 2)$.

2.4 Indexing or Blocking Techniques:

Traditional ER requires pairwise comparisons between all the records. For instance, gievn two entity collections E_1 and E_2, with sizes $|E_1|$ and $|E_2|$, it requires $|E_1| \times |E_2|$ comparisons between entity records. In practice, this is infeasible when, $|E_1|$ and $|E_2|$ are large due to the inherent quadratic complexity of the comparison process.

Indexing or blocking reduces the number of detail-level pairwise comparisons between records by removing pairs that are unlikely to be real matches. The traditional blocking techniques partition the databases into non-overlapping blocks, only comparing the records within blocks. Hence, reducing the number of pairwise comparisons.

Our method is an example of join-based blocking techniques that convert blocking into the nearest neighbour search. The blocks are created by searching the vector space for similar records using the Em-K indexing rather than partitioning the dataset. As a result, similar records are grouped into overlapping blocks by combining spatial joins with block building [26].

3 PROBLEM FORMULATION

In this section, we formulate our problem using the concepts and the definitions presented in the previous section. The proposed Em-K indexing method functions as a preprocessing approach for a more detailed query-matching ER. We use deduplication as a motivating example that explains the basic building blocks to the query matching problem. Thus we defined the following two problems:

Query Matching: is the problem of finding similar records given two entity collections, E_r a reference database, and Q a stream of queries. The size of each dataset is denoted by $|E_r|$ and $|Q|$ respectively and we assume $|E_r|$ is fixed and $|Q| \to \infty$. Let $Q = q_1, q_2, .., q_{|Q|}$, where each query q_i in Q represents a record of an entity e_j. A query record q_i has the same attribute schema as the records in E_r. Hence for each query q_i in Q, the records in E_r that belong to the same entity (e_j) need to be found. In real-life problems, the query rate might be very fast. There may or may not be a matching record for every $q_i \in Q$ that belongs to the entity e_j.

PROBLEM STATEMENT 1. *(Indexing for Query matching): For two datasets (one is a reference database and the other is a stream of queries) with overlapping records, run the best algorithm for a given amount of time to find the block of records containing as many matches as possible for the querying record. This is a pre-processing step to increase the efficiency of the subsequent detailed query match.*

We also consider the deduplication problem here, primarily as an explanatory model. We defined deduplication in Section 2.2 under *dirty ER*. In the following, we define the indexing for deduplication.

PROBLEM STATEMENT 2. *(Indexing for deduplication): For a given entity collection containing duplicates, group similar records into blocks to reduce the number of comparisons needed in subsequent detailed deduplication while missing as few matches as possible.*

Traditional indexing techniques split the database into non-overlapping blocks, only comparing the records within any block [8, 11, 16, 25, 26]. Blocks of similar pairs are determined by building an indexing structure that takes a set of records as input and classifies them according to some criteria. Usually, this criterion is based on matching a *blocking key* consisting of a single or several attribute values of records [8]. Our method uses blocking values to create blocks of records that transform blocking into a k-nearest neighbour (k-NN) search. We combine the spatial joins with block building to convert blocking values to a similarity preserving Euclidean space. The result is overlapping blocks of records.

The proposed method requires mapping blocking values into multidimensional vectors. Since many comparisons in ER are between strings of characters, we focus on entity attributes that contain string values here. Unlike the similarity between string values, the similarity between numerical values is easy to compute using L^p norms in a metric-space [18]. Hence, we used this property of the metric-space to propose a scalable indexing technique for query matching.

If we assume strings as elements of a complicated high-dimensional space, the distance between two different strings is typically large. However, misspelled strings tend to be located near correctly spelled strings [24]. The embedding of a string database into a metric-space needs to preserve these two properties. Thus, coordinates for blocking values are determined in a Euclidean space such that the associated Euclidean distances approximate the dissimilarities between the original blocking values.

The general problem of assigning coordinates in this manner is one of embedding a metric or non-metric-space into a Euclidean space [31]. Suppose R is a collection of objects, δ measures the distances between R objects, X represents the coordinates matrix for the R objects in the Euclidean space, and d measures the distances between coordinates.

Embedding of a metric or non-metric-space (R, δ) into a Euclidean space (X, d) is a mapping $\phi : R \rightarrow X$. In this paper (R, δ) will always be a finite space (*i.e.*, R is a finite set) and (X, d) will always be a Euclidean space.

PROBLEM STATEMENT 3. *(The embedding problem): For a given metric or non-metric space, find a ϕ that minimizes distortion, stress, or a similar error metric between (R, δ) and (X, d).*

A commonly used technique for embedding a set of distances (or dissimilarities) into a Euclidean space is multidimensional scaling (MDS) [19]. We applied MDS because we can adapt it to achieve good time efficiency and distance preserve capability for large-scale data with a small amount of extra effort.

Among the variants of MDS, we use least-squares multidimensional scaling (LSMDS) for the embedding since it gave the best results compared to other variations such as classical scaling [14]. We can map a set of blocking values to a lower-dimensional Euclidean space by applying LSMDS such that the distances between vectors preserve the dissimilarities between them. This embedding leaves similar blocking values closer in the Euclidean space allowing efficient, geometric-based indexing.

LSMDS initially maps each item in the non-metric or metric-space to a K-dimensional point. Then minimises the discrepancy between the actual dissimilarities and the estimated distances in the K-dimensional space by optimisation [13]. This discrepancy is measured using *raw stress* (σ_{raw}) given by the relative error where δ_{ij} is the dissimilarity between the two objects and d_{ij} is the Euclidean distance between their estimated points.

$$\sigma_{raw}(\mathbf{X}) = \sum_{i,j=1}^{n} w_{ij} \Big(d_{ij}(\mathbf{X}) - \delta_{ij} \Big)^2. \tag{1}$$

Possible weights for each pair of points are denoted by w_{ij}. Weights are useful in handling missing values and the default values are $w_{ij} = 0$, if δ_{ij} is missing and $w_{ij} = 1$, otherwise [13]. We do not apply weights in this work, hence, $w_{ij} = 1$ always. We prefer the normalized stress (σ) in our experiments since it is popular and theoretically justified. The normalized stress (σ) is obtained by $\sigma = \sqrt{\sigma_{raw}(\mathbf{X})/\delta_{ij}^2}$.

However, traditional MDS algorithms such as LSMDS require extensive preprocessing and usually are computationally expensive, thus not appropriate for large scale applications. The two main drawbacks are,

- MDS requires $O(N^2)$ time, where N is the number of items. Thus, it is impractical for large N.
- In an out-of-sample setting or a query-by-example setting, a query item has to be mapped to a point in the pre-mapped Euclidean space. Given LSMDS algorithm is $O(N^2)$, an incremental algorithm to search/add a new item in the database would be $O(N)$. Hence a query search would be similar to sequential scanning of a database [12].

Among the proposed methods of scalable MDS, we are interested in using an out-of-sample embedding approach as a scaling method for LSMDS. We have two main purposes:

(1) To embed large-scale reference databases.
(2) To embed previously unseen data to a pre-mapped Euclidean space.

Suppose we have a configuration of N points in a K-dimensional Euclidean space obtained by applying LSMDS to a set of N objects. Let Q be out-of-sample objects, with measured pairwise dissimilarities from each of the original N objects. The out-of-sample embedding problem is to embed the new Q objects into the pre-mapped K-dimensional Euclidean space.

Our out-of-sample embedding approach uses the stochastic gradient descent algorithm to minimise the following objective function for numerical optimisation. The out-of-sample embedding of a new object y is obtained by minimising the following objective function,

$$\hat{\sigma}(y) = \sum_{i=1}^{n} \big(\|x_i - \hat{y}\|_2 - \delta_{iy} \big)^2, \qquad (2)$$

where y is the new object and $\hat{y}$ is its position in the Euclidean space. The δ_{iy} represent the dissimilarities between point i and the new object y. The Euclidean distance between the i^{th} point and y in the Euclidean space is given by $\|x_i - \hat{y}\|_2$. We seek to find a position of y that minimises $\hat{\sigma}(y)$. Here we keep $w_{ij} = 1$, similar to LSMDS.

The out-of-sample embedding approach becomes inefficient for large-scale data by comparing each new point with all the existing points. We scale our out-of-sample embedding solution to accommodate large-scale data by only considering a fraction of the pre-mapped Euclidean space. The initially selected subset of pre-mapped data is usually known as *landmarks*. Landmarks or *anchors* have been used with out-of-sample extensions to scale MDS and other embedding techniques [28, 31]. We discuss the characteristics of a good set of landmarks and landmarks selection methods in Section 5.

Once the Euclidean space consists of all the data required, we then formulate our indexing method to generate blocks that group close-by vectors in the Euclidean space.

MDS maps high dimensional data (in the original space) into a Euclidean space (vector space). The rationale for performing such a mapping is to approximate the distances between objects in an original space in Euclidean space. Searching for similar points in Euclidean space is less expensive and quick since we can use efficient data structures such as Kd-trees. We use Kd-trees and k-NN search to find similar vectors in the Euclidean space.

In general, k-NN search refers to finding the closest elements for a query q within a given set of points N, as measured by some distance function $d(N, q)$ [20]. The distance function d is a metric, e.g., L^p norm, which satisfies the non-negativity, identity, symmetry, and triangle inequality properties [4, 23]. We use Euclidean distances $d_E($ where $p = 2)$ in our calculations. Here we are interested in finding the k-NNs (k nearest neighbours) where k may be moderately large.

PROBLEM STATEMENT 4. *k-nearest neighbour (k-NN) search: The query retrieves the k closest elements to q in N. If the collection to be searched consists of fewer than k objects, the query returns the whole database. If the set of the k nearest neighbours of q are n_c, then formally, n_c can be defined as follows:*
$$kNN(q) = n_c \subseteq N, |n_c| = k \wedge \forall x \in n_c, y \in N - n_c : d(q, x) \leq d(q, y).$$

An index is a data structure that reduces the number of distance evaluations needed at query time. An efficient and scalable indexing method can facilitate accurate and efficient k-NN search that supports large-scale datasets. We can apply several k-NN search methods for indexing arbitrary metric spaces; for more details, refer to the surveys [5, 15]. Distance-based indexing methods use distance computations to build the index. Once the index is created, these can often perform similarity queries with a significantly lower number of distance computations than a sequential scan of the entire dataset [15].

The decision of which indexing structure to apply depends on several factors, including query type, data type, complexity and the application. Among many data structures, we choose Kd-trees for the k-NN search. It is considered one of the best data structures for indexing multidimensional spaces and is designed for efficient k-NN search [1].

Kd-trees organise K-dimensional vectors of numeric data. Each internal node of the tree represents a branching decision in terms of a single attribute's value, called a *split value* [30]. These internal nodes generate a splitting hyperplane that divides the space into subspaces using this split value, usually the median value along the splitting dimension. We will use the median when constructing the Kd-tree for the data in our experiments. Building a Kd-tree (with the number of dimensions K fixed, and dataset size N) has $O(N \log N)$ complexity [3]. For more details on Kd-trees and k-NN search implementations, refer to [3].

The k-NN search algorithm aims to find a node in the tree closest to a given input vector. Searching a Kd-tree for k nearest neighbours is $O(k \log N)$, which is the key to fast indexing. It uses tree properties to quickly eliminate large portions of the search space [1].

4 METHODS OF EM-K INDEXING

4.1 Indexing for Deduplication

Deduplication refers to identifying matching records within a single database and has many applications in database and business contexts. For instance, many businesses maintain databases of customer information that are utilised for advertising purposes, e.g., emailing flyers. Duplicate entries might arise because of errors in data entry or address changes. Deduplication techniques are useful to remove duplicate entries and to improve the quality of the collected information. Duplicate-free customer information databases prevent emailing several copies of flyers to the same customer, which reduces the cost of advertising, but there are many other benefits.

Indexing is a preprocessing step to avoid the need to perform detail-level comparisons between $O(N^2)$ pairs of records. The proposed indexing method utilises the properties of vectors and Kd-trees in a Euclidean space. The method has two main steps,

(1) Embedding the blocking values: We select the blocking criteria based on the attributes of the records in a dataset. For instance, given a set of records with entity identifying attributes

such as first name, last name, date of birth, or postcode, we may choose one or several values of them in our indexing method. We embed the blocking values of a dataset/database in the Euclidean space by applying LSMDS. The embedding depends on the size of the dataset. We propose two techniques:

(a) *Complete LSMDS*: For a given dataset of size N, apply LSMDS for the blocking values of all the records.

(b) *Landmark LSMDS*: Apply complete LSMDS only to a fraction of the dataset (the landmark records). Then the remaining records are embedded using the out-of-sample embedding against the landmark points. We explained this approach in Section 2.

(2) Nearest neighbour search: Searching for similar points is a two-step procedure.

(a) The first step is to build the Kd-tree in the Euclidean space using all the points that represent the blocking values in this space.

(b) The second step is to create blocks of similar points by searching the nearest-neighbours of the Kd-tree nodes.

Since the Kd-tree construction uses all the available points in the Euclidean space, each record becomes a node in the Kd-tree. Likewise, each node becomes a query against the rest of the nodes in the k-NN search. Each node has a fixed number of nearest neighbours (NNs) allocated for them as we keep the k-NN search fixed for every querying node. Hence, each node in the tree becomes a small block of records that contains its NNs as the members. The block sizes are uniform and depend on the number of NNs allocated for a node.

Once all the blocks are determined by the indexing method, we return the pairs of similar points in each block as potential matching records. Then, we retrieve the original blocking values that correspond to these points and compare them to classify the pairs as candidate matches or non-matches based on a pre-selected threshold.

A detailed level comparison among the other attribute values will be only required between the pairs that indexing identifies as candidate matches. Thus our indexing method act as a filtering step that reduces the total number of detailed comparisons one has to perform when identifying similar records in a dataset.

4.1.1 Complexity. We can quantify the complexity of the proposed indexing method. The method has two components: a relatively slow step where the records (blocking values) are mapped to a Euclidean space, followed by a relatively fast step that creates blocks in the Euclidean space.

Assume that we have N records in the underlying database. Complete LSMDS requires calculating the distance between all pairs of blocking values, hence, $O(N^2)$ operations for the embedding. This complexity dominates the LSMDS calculations. However, for large N, we can choose a set of landmarks L and apply LSMDS with a complexity of $O(L^2)$. The embedding of the rest of the points has a linear complexity of $O(ML)$ operations, where $M = N - L$. Hence the overall complexity of landmark LSMDS is $O(L^2 + ML)$.

The second phase builds and searches the Kd-tree to create blocks of records. Building the Kd-tree requires $O(N \log N)$ operations, and searching for k nearest neighbours for N points requires $O(Nk \log N)$ operations, where the size of a block (B) is equal to k. Hence the overall complexity of the indexing step is $O((1 + k)N \log N)$.

Since a complete LSMDS requires $O(N^2)$ operations, we recommend using the landmark LSMDS and as small as possible k to reduce the complexity of the proposed method.

4.2 Indexing for Query Matching

Query matching in this paper refers to querying a stream of records against a reference database to find records that refer to the same entity as the query (see Section 3). Query processing should be

quick and accurate for many ER solutions to increase their usability in real-time applications. In some applications, a stream of queries might need to be processed within a given time, collecting as many matches as possible. In such settings, we have to trade accuracy against speed when detecting matching records. This section presents a scalable indexing method for real-time, approximate query matching against a large-scale database. The ideas presented in Section 4.1 serve as the basic building blocks for the proposed method.

Following the *clean-clean ER* scenario, we consider a large-scale reference database E_r and a batch of streaming queries Q. Similar to the previous indexing method, we first embed all the blocking values of the reference database in a Euclidean space. We use a set of landmarks and the out-of-sample embedding of LSMDS to reduce the overhead of embedding the database E_r. The embedding is a two-step procedure: 1) apply LSMDS over the landmarks 2) then map the rest of the blocking values using out-of-sample embedding of LSMDS based on the distances to these landmarks. We then construct a Kd-tree using all the points in the Euclidean space, where each data point becomes a node in the tree.

For streaming queries, we process a single query record at a time. First, we embed the blocking value of the query record in the pre-mapped Euclidean space. In general, out-of-sample embedding would require calculating all the distances from the new query to the pre-mapped blocking values in the original string space, which is not desirable with large-scale databases. Hence, we only calculate the distances to the landmarks when mapping a new record. The mapping position is determined by applying the out-of-sample embedding of LSMDS to the new query record. The process is similar to mapping the blocking values that are not landmarks in the reference database.

The only inputs we need for the mapping are the distances from a query to the pre-selected landmark blocking values. Then we search for the k-NNs of the new point using the existing Kd-tree structure of the reference database. A new block of points that contains similar points for the query point will be determined. The block size B depends on the k of the k-NN search. We retrieve the original blocking values for the block of similar points and use a pre-selected threshold to filter out the potential matching pairs. The records that contain these as blocking values are the matching records to the querying record. A detailed level comparison between the candidate matching pairs may be required after the initial step of indexing.

The algorithm needs to process a query within a fixed time, potentially sub-second. Since we are querying a stream of query records against a large-scale database, the processing time of a single query is limited. The overall system performance can be improved by trading accuracy against efficiency and scalability. We will fine-tune our method to trade-off many comparisons against accuracy to offer a greater number of detections within a fixed time. Thus we select a set of optimal parameters that scale our data by considering the trade-off between accuracy and scalability.

4.2.1 Complexity. We can quantify the complexity of different stages of the Em-K indexing for query matching. Similar to the previous method, it contains two components: a relatively slow step where a new query is mapped to a pre-mapped Euclidean space, followed by a relatively fast step that creates a block of similar points which contains potential matches for the query. We assume below that the reference database has already been embedded since this cost is amortised across many queries.

Suppose that we have N data points in our reference database. Typically it would require $O(N)$ operations to compare the existing records with a query point, which is not feasible for larger N. We avoid this complexity by applying out-of-sample embedding of LSMDS using a set of L landmarks, which requires only $O(L)$ operations to embed a new data point. The embedding is efficient if L is chosen such that $L \ll N$.

In the second phase, we search for k nearest neighbours for the query. Therefore, the block size B is equal to k. It will cost $O(k \log N)$ operations to search the tree and $O(k)$ operations to compare a new query point within a block. The total cost of indexing is $O(L + k \log N)$. However, $k \ll L \ll N$ in large-scale applications. The cost of the k-NN search is insignificant due to the efficient indexing structure of the Kd-tree. Thus out-of-sample embedding step dominates the complexity of the proposed method.

The number of landmarks L will affect both the accuracy and scalability of this method. We will discuss the selection of landmarks and results in the next section.

5 EXPERIMENTAL VALIDATION

We evaluated the two proposed Em-K indexing methods under various settings (e.g., different dimensions, varying block sizes and datasets, and error rates). Two main questions to study in the experiments are: (1) Are these proposed methods robust over various settings? (2) Does Em-K indexing achieve high accuracy and good scalability?

All algorithms are implemented in R and executed on a desktop with Intel Core 5 Quad 2.3GHz, 16GB RAM, and MacOS Big Sur.

5.1 Set Up

5.1.1 Data Sets. We examined the performance of our methods over two synthetic datasets. They can be manipulated to have significant variations in their size and characteristics (e.g., error rates).

- **Dataset-1:** The first data set contains records with synthetically generated biographic information. Each record has a given name and a surname. They are generated using the tool Geco [10]. We introduced duplicate records with errors by slightly modifying the values of randomly selected entries. In record generation, we assumed a record has only one duplicate with a maximum of two typographical errors (substitutions, deletions, insertions, and transpositions) in both attribute values.
  For the deduplication datasets, there is one duplicate for a particular record within the dataset. Similarly, in query matching, each query has one duplicate within the reference database, and the reference database is duplicate free.
- **Dataset-2:** The second dataset is the benchmark dataset presented in [27]. It is based on personal records from the North Carolina voter registry and synthetically generated duplicates using Geco [10]. Each record has several attribute fields. We cannot control the errors of the duplicate records in this dataset since it has been formulated as a benchmark dataset. However, after careful analysis of the dataset, we estimated that a duplicate record has a maximum of three edit distance errors for this work. We have only considered the first name and the last name fields in our experiments.
  Similarly, we selected the deduplication data such that there is only one duplicate for a particular record within the dataset. In the query matching, we choose the queries to have only one duplicate within the reference database and the reference database to be duplicate free.

5.1.2 Matching Rates. We control the number of duplicates in our experiments in order to understand how well the method works in different circumstances. We used two matching rates: Consider the datasets E, E_r and Q with sizes $|E|$, $|E_r|$ and $|Q|$,

(1) *Deduplication matching rate (DMR)*: the matching rate for deduplication is defined to be the $DMR(E) = \frac{E_d}{|E|}$ where E_d is the number of duplicate records in E.

(2) *Query matching rate (QMR)*: the matching rate for query matching of Q against a reference database E_r is defined to be the $QMR = \frac{M_{RQ}}{|E_r|}$ where M_{RQ} is the number of records in E_r that has a matching record (duplicate) in Q.

We control the number of duplicates within and between datasets by changing the DMR and QMR. We represent DMR and QMR as percentages in our experiments.

5.1.3 Performance Evaluation Metrics. We use two measures to quantify the efficiency and the quality of our indexing method proposed for deduplication [9].

- Reduction Ratio (RR): Measures the relative reduction of the comparison space, given by $RR = 1 - \frac{N_b}{|E|(|E|-1)}$ where N_b is the number of potential matching pairs produced by an indexing algorithm. This quantifies how useful the indexing is at reducing the search space for detailed comparisons.
- Pair Completeness (PC): This is given by $PC = \frac{N_m}{M}$, where N_m denotes the detected number of real matches by the indexing algorithm and M represents the number of all real matches in E. This is a measure of how accurate the indexing is.

Both PC and RR are defined in the interval $[0, 1]$, with higher values indicating higher recall and efficiency, respectively. However, PC and RR have a trade-off: more comparisons (higher N_b) allow high PC but reduce the RR. Therefore, indexing techniques are successful when they achieve a fair balance between PC and RR.

We used two measures to evaluate the performance of our indexing method of query matching. These are different from the standard measures of indexing defined above as we combine indexing with query matching here. Hence, in this context, we are interested in measuring the efficiency of the method in terms of time and speed of processing a query.

The Em-K indexing method returns a block of records that contains both true positive (TP) and false positive (FP) matches per query as a final result. Hence the following measures are used.

- Number of true positives per computational effort: Measures the number of true matching records determined by the indexing method when processing queries within a set period.
- Precision: Measures the accuracy of the query matching in terms of precision. The precision P is denoted by $P = \frac{|TP|}{|TP|+|FP|}$, where $|TP|$ is the number of true positives and $|FP|$ is the number of false positives.

All the CPU running times are measured in *seconds* and denoted by RT.

5.2 Choice of Parameters

Several factors may impact the performance of the proposed methods; some of them (e.g., dimension (K), block size (B), landmarks (L)) are control parameters that we rely on to fine-tune the performance, while others (e.g., dataset size or error rate) are parameters determined by data sets.

In this section, We initially discuss the choices of the K, B, L parameters in our indexing implementations (and their rationale). The robustness of the methods is measured with respect to the varying sizes and matching rates of the data sets in Sections 5.5–5.7.

K: The dimension of the Euclidean space (K) plays an important role in the performance of our indexing methods. We applied LSMDS to a sample of 5000 records selected from Dataset-1. In Figure 1, the first y-axis shows the stress (σ defined in Equation 1) decreases as K increases.

A good value of K should differentiate similar objects from dissimilar ones by approximating the original distances. If K is too small, we will have high-stress values where dissimilar pairs will not fall far enough from each other. It could also place dissimilar pairs closer and similar pairs further apart. Conversely, high K values will have low-stress values. However, in terms of RT, higher

dimensions increase the embedding time. In Figure 1, the second y-axis represents embedding RT. It takes more than 30 minutes to embed the dataset in 19 or 20 dimensions.

Considering the trade-off between the *stress vs dimension* and the *embedding time vs dimension*, a reasonable value of σ is found around 6-8 dimensions. This is consistent with the embedding name strings in the Euclidean space, as discussed in detail in [14]. We will use $K = 7$ here.
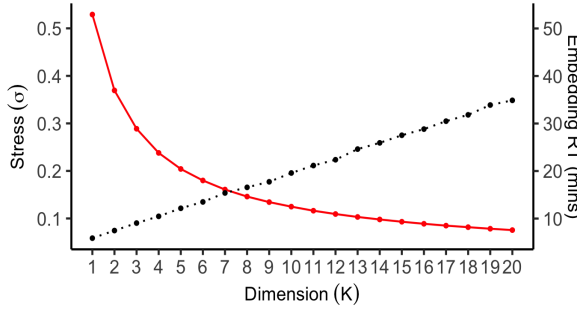

Fig. 1. The trade-off between the dimension (K) vs stress (σ) and embedding time. The first y-axis represents σ, while the second y-axis represents the embedding time. The σ tends towards a small but non-zero asymptote when K increases. The running time increases linearly when K increases. Higher dimensions allow lower σ values for the embedding but increase the embedding time, for marginal benefit.

B: Block size is a dominating factor that directly affects the effectiveness and efficiency of many indexing techniques. Large block sizes increase RT in the indexing step and have a low RR and high PC values. In contrast, small block sizes lead to high RR values with fewer comparisons within each block. However, this may result in low PC values due to missing some matches. Blocks of similar points are determined by k-NN search in Em-K indexing methods. Hence, B is equal to k (number of nearest neighbours). We will consider the choice of B in detail in Section 5.2.1.

L: The number of landmarks (L) is another important factor in our indexing methods. Landmarks are utilised for two different purposes in this work. First, we use landmarks for embedding large-scale data into a Euclidean space when the standard LSMDS method becomes inefficient. Second, we use landmarks to embed the out-of-sample queries in a pre-mapped Euclidean space. We discuss the role of landmarks in deduplication and their impact on the proposed indexing method in Section 5.2.2.

In the following experiments we investigated how B and L impact our indexing algorithms.

5.2.1 Varying Block Sizes (B). To test the choice of B, we used sample datasets containing 5000 records from the two data sources. We set the $DMR = 10\%$ for the data selected from Dataset-1, which means there are 500 duplicates within the selected 5000 records in the sample dataset. Similarly, we used $DMR = 7.5\%$ for the second data sample selected from Dataset-2. Hence within the 5000 records, 375 of them are duplicates.

We performed deduplication indexing on the two datasets. Figure 2 illustrates the trade-off between PC and RR of our indexing method for different block sizes using the Dataset-1. In each instance, we changed B by varying k in the k-NN search. PC increases with the increase of B. In contrast, RR decreases due to the increment in the number of comparisons within each block of records.

The results also indicate that dimension around 7 are good at shifting the PC-RR curve to the top-right corner delivering a good ratio between RR and PC. However, higher dimensions also
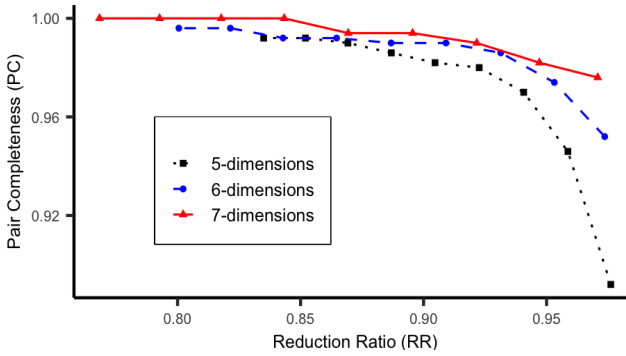
Fig. 2. The trade-off between the reduction ratio (RR) and pair completeness (PC) in different dimensions. Ideally, RR=PC=1, hence we prefer methods whose results lie as close to the top-right corner of the graph. The block sizes are 100, 90, 80, 70, 60, 50, 40, 30, 20. Large blocks achieve higher PC but lower RR. The three curves illustrate that the higher the dimension, the better the results, up to the point of diminishing returns.

mean higher computational costs (e.g., high RT) for the embedding. Based on Figure 1 and Figure 2, therefore, we conclude that using $K = 7$ and $B = \{50, 60\}$ gives the best compromise PC-RR ratio and RT overall for the given data set. In subsequent experiments, we used $K = 7$.
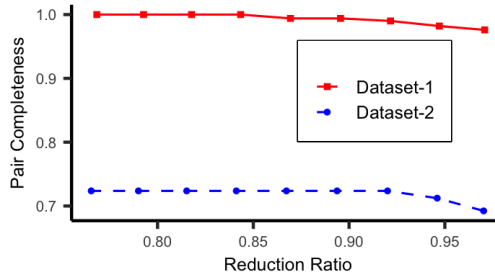

Fig. 3. The trade-off between the reduction ratio (RR) and pair completeness (PC) for two different datasets for the proposed indexing. The block sizes are 100, 90, 80, 70, 60, 50, 40, 30, 20. PC is very low for the second dataset while RR values are closer.

In most comparisons, we observed similar results for both datasets and discussed only the results of Dataset-1. However, we present one comparison that has comparably different results here. Figure 3 compares the two datasets in a fixed dimension ($K = 7$), varying B in each instance. Both data sets achieve similar RR values with different PC values. PC values are comparatively low for Dataset-2 and are around 70% for most occurrences.

Dataset 1 shows better results compared to Dataset-2 in Figure 3. This behaviour is expected since the two datasets have different characteristics, e.g., different matching rates and a different number of errors in each field. Furthermore, we used different pre-selected thresholds (θ_m) when validating candidate matching pairs in the two datasets. These thresholds are selected based on the errors in the two datasets. In Dataset-1, duplicates have a maximum of two typographical errors and therefore $\theta_m = 2$. For Dataset-2, we assumed that each duplicate record has a maximum of three typographical errors, and we set $\theta_m = 3$ therein.

5.2.2 The Effect of Landmarks. The following experiment investigated the effect of the two different embedding techniques on the proposed indexing method. Performance is measured using the PC

and RR curves. First, we applied complete LSMDS similar to Figure 2, keeping the $K = 7$ fixed. Second, we applied LSMDS only to a set of landmarks in the same dimension. The remaining points are embedded using the out-of-sample embedding of LSMDS and the distances to landmarks. Blocksize B is varied as before.
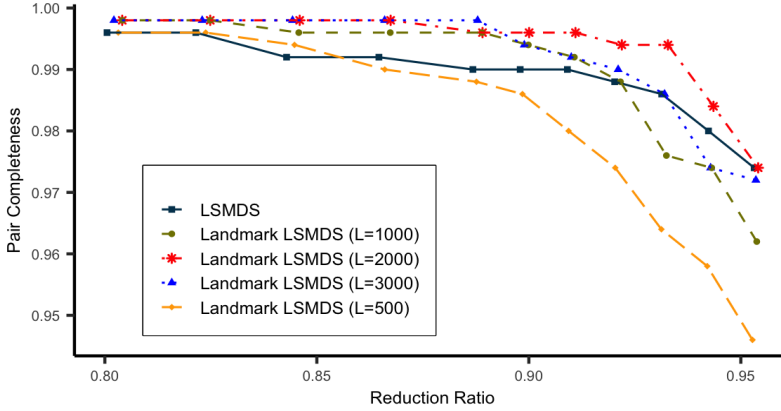

Fig. 4. The trade-off between the reduction ratio (RR) and pair completeness (PC) for the proposed indexing based on different embeddings.The results are based on a complete LSMDS and landmark based LSMDS for different number of landmarks. The block sizes are 100, 90, 80, 70, 60, 55, 50, 45, 40, 35, 30.

Figure 4 compares the trade-off between PC and RR for complete LSMDS and landmark LSMDS (for different L). Each instance represents a different block size. PC and RR change similarly to the previous experiment for different block sizes. However, Figure 4 suggests that we can get similar results by choosing an approximate embedding that uses landmarks instead of complete LSMDS. The use of landmarks decreases the distance calculations and we can avoid the inherent complexity and inefficiency of LSMDS when processing large-scale data.

Using our indexing solution, we can solve deduplication applications in ER. The method requires embedding records into a Euclidean space in order to apply the indexing technique. Since complete LSMDS is not suitable for large-scale data, we recommend using the landmark LSMDS. We applied the farthest first sampling [17] for reproducible results in landmarks selection; however, random selection works well in practice.

The optimal parameter setting for our data is $K = 7$, $B = 50$ and $L = 1500$. We used the proposed indexing method of deduplication as a motivating example for the next set of experiments.

5.3 Indexing for Query Matching

In our experiments, we used a reference database with $E_r = 5000$ records. The streaming query dataset Q is flexible in size because we only consider streaming queries within a period. Each query in Q has a duplicate record in the reference database, i.e., $QMR = 1$. We made this assumption to keep the experiment more efficient instead of mimicking a real-world ER problem. In a real-world scenario, each query may not have a matching record within the reference database, or the same query may appear in the stream to search against the reference database. However, the time required for searching is the same when no match is present.

Similar to the previous method, several factors affect the performance of the proposed indexing method for querying, e.g., K, L, B. Since the embedding of the reference database is the same as before, we keep $K = 7$. The block size is B is equal to k, the number of nearest neighbours in k-NN search.

We used landmarks to support the out-of-sample embedding. The number of landmarks, L, directly impacts the running time (RT) of the proposed method since each query needs to be embedded in the Euclidean space. The other costs that contribute to RT are the distance calculations and k-NN search.

First, we embedded the reference database generated from Dataset-1 applying landmark LSMDS. We chose the landmarks based on the farthest-first sampling. Then, we built a Kd-tree using all the reference data points in the Euclidean space. Once the Kd-tree is built, we passed queries to search the tree for k-NNs. Hence, each query in Q needs to be mapped in the Euclidean space. We used the same set of landmarks among the reference data points to map the query points. Distance calculations are required among the new query point and the landmarks when applying the out-of-sample embedding of LSMDS. A block of similar points is determined for a new query by searching k-NNs in the existing Kd-tree.

Our method process a single query at a time. In the following experiment, we processed 500 queries against the reference database. For each query, we measured the embedding RT and the distance calculation RT separately. Then we calculated the mean values for both categories, processing all the 500 queries. Figure 5 shows the comparison of the mean RT of distance calculations and out-of-sample embedding for varying numbers of landmarks. Increasing L linearly increases the embedding RT of a single query. The distance calculation RT also increases linearly with L but is negligible compared to out-of-sample embedding RT.
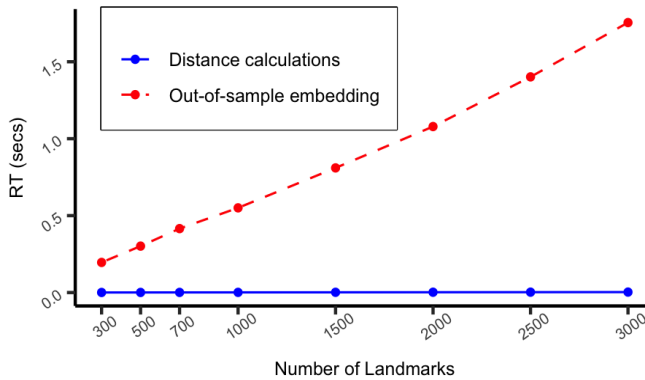

Fig. 5. The calculation times for distance calculations and out-of-sample embedding for query matching. Both depend on the number of landmarks, but distance calculations are much faster.

We also measured the cost of the k-NN search when creating a block of records for a new query. This search can be done efficiently in the Euclidean space using the Kd-tree and priority queues. It takes less than a millisecond, which is insignificant compared to the total RT of the embedding process. Moreover, increasing k has a smaller impact due to its efficient implementation with priority queues [3].

A scalable query matching method should be able to process as many queries as possible within a period. In our indexing method, increasing L limits the number of queries processed within a set period. On the other hand, small L tends to decrease the accuracy of the embedding. As a result, a new query may be not mapped closer to its duplicate, reducing the probability of grouping them as similar points. Hence an optimal L is required to maintain the scalability without degrading the quality of the results.

We measured the scalability and the efficiency of our method using two quantitative measures with respect to time: the number of true positive ($|TP|$) matches detected per computational effort and the precision (P) per computational effort. Hence we processed a stream of queries Q against a reference database E_r within a given period, varying the control parameters such as L and B in different instances. Then we calculated $|TP|$ and $|FP|$ found by our method in each instance. The accuracy of the results is measured in terms of precision. It measures the rate of TP against all the positive results (sum of $|TP|$ and $|FP|$) returned by the method within the given period. Subsequently, we determined an optimal set of parameter values for our data that returns the highest TP matches and precision within a fixed period.

In the following experiments, we used a reference database of E_r = 5000 records and the stream of query records Q = 500. We applied landmark LSMDS to embed the records in E_r to the Euclidean space in each experiment. Then the queries are processed using the same set of landmarks. Hence, every instance of a different L has a disparate embedding of the reference database, then used for query embedding and searching.
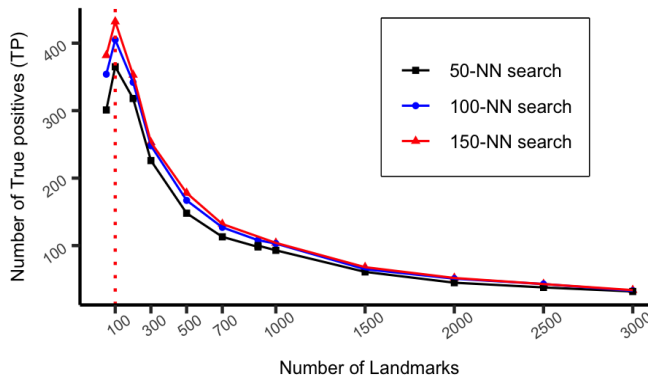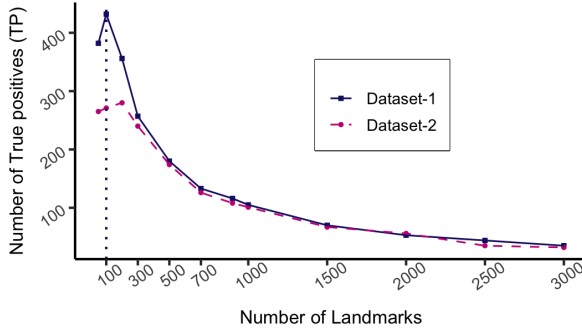

Fig. 6. The trade-off between the number of true positives and the number of landmarks for three different block sizes. Here the block size is equal to the number of k-NNs. Many landmarks allow fewer queries to be processed, while many k-NNs allow more true match detections.

Figure 6 shows the comparison of the $|TP|$ against varying the number of landmarks and k-NNs. In each instance, we processed queries within a fixed period (60 seconds). Increasing L decreases $|TP|$ because more landmarks allow fewer queries to be processed. This phenomenon is expected since large L increase the running time of processing a single query. With more landmarks, there is a higher probability of finding matches for those queries due to the increasing accuracy of the embedding. In contrast, few landmarks allow more queries to be processed within a period since the running time (RT) for embedding a single query is small.

Figure 6 suggests that we only need 100 landmarks to detect the highest number of TP matches for the given data. The method has processed all the 500 queries within a minute, detecting 432 TP matches. The average time for processing a single query is 0.07 seconds.

Based on Figure 6, we conclude that setting k = 150 and L = 100 gives the best trade-off between the quality of the results and the RT. This result is consistent with existing real-time query matching techniques that process a query within a sub-second time [22].

To validate the robustness of the method, we compared the results of applying the EM-K indexing method to two reference databases and stream of queries derived from Dataset-1 and Dataset-2, respectively in Figure 7. Both reference databases contained 5000 records. The two streams of queries

(a) The trade-off between the number of TP and the number of landmarks for two different datasets.
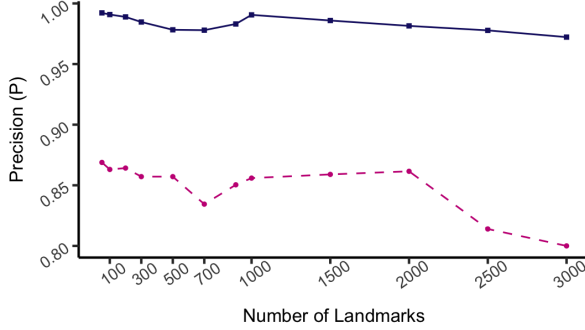
(b) The trade-off between the precision and the number landmarks for two different datasets.

Fig. 7. The two figures (a) and (b) compares the Dataset-1 and Dataset-2 in terms of number of TP and Precision varying the L The k-NNs are fixed to $k = 150$ and $T = 60$ seconds in each instance. The curves illustrate similar trends for similar parameter settings in (a). However, in (c), Dataset-2 exhibits low precision compared to Dataset-1.

are processed against the two reference databases separately with similar parameter settings. The query matching rate (QMR) is equal to 1 for both. Keeping the control parameters fixed at $K = 7$, $B = 150$ and $T = 60$ seconds, we vary L to compare different results.

Figure 7-(a) compares the trade-off between the $|TP|$ per computational effort, and the number of landmarks. Increasing L decreases the $|TP|$ for query matching in both datasets similar to Figure 6. Few landmarks allow processing all the queries within 60 seconds, whereas many landmarks allow processing fewer queries. While both reference databases are similar in size, the total queries processed are different in size. A total of 500 queries are processed against Dataset-1, and 375 queries are processed against Dataset-2 in the experiment. Hence, we observe fewer TP matches for Dataset-2 than Dataset-1 in their highest performance. However, the curves illustrate similar trends for both datasets.

Based on the results produced by the previous experiment, Figure 7-(b) compares the two datasets in terms of the precision (P) against varying L In each instance, the queries are processed within T=60 seconds. Hence the P values are measured per computational effort. Dataset-2 has lower P values compared to Dataset-1. This behaviour is expected since we assume that the duplicate records in Dataset-2 contain a maximum of three edit distance errors. The pre-selected thresholds (θ_m)

are different for the two datasets. However, we do not have the exact details of error rates for this benchmark dataset. Hence, the method may allow more false positives in the final block of records retrieved for a given query, reducing the P.

However, the overall results suggest that the proposed method is robust over different datasets. The optimal parameter set that provides the best compromise for our data is $K = 7$, B (or k)= 150 and $L = 100, 300$. We need a few landmarks to achieve the most number of TP matches for a stream of queries within a period. Few landmarks make the embedding efficient. As a result, a single query can be processed within less than a second against a reference database to find matches. Hence our method scales well for approximate query matching against a large-scale reference database for ER.

6 DISCUSSION

The Em-K indexing methods embed a set of strings in a metric-space, particularly a lower-dimensional Euclidean space. There is a trade-off between the dimension and the accuracy of the embedding. Higher dimensions allow the embedding to be more accurate. However, it does not scale well for large datasets. In our experiments, we only selected two blocking variables for embedding and indexing. We can also use other blocking variables, such as addresses and gender. As a result, the dimension of the Euclidean space needs to extend accordingly to facilitate those.

The most costly part of the method is the amount of time to embed queries, which increases linearly with the number of landmarks. However, this approach is easily parallelizable since each query is processed separately. We can tune the parameters for a fast, less accurate query matching or a slow, more accurate method depending on the application.

Our method is designed to solve approximate query matching rather than exact query matching. It means we expect the queries to contain errors in their attribute values. We could easily perform an exact search based on lexicographical order considering the query and the reference dataset using a data structure such as a binary tree to find exact matches as a pre-filter for our method.

Our indexing method for query matching uses distance computations to embed the reference database in the Euclidean space. This embedding is a slow process that requires a minimum of $O(L^2)$ operations where L is the number of landmarks. We consider this as the training phase, which only needs to be performed once. However, with the Kd-tree built, we can perform similarity queries with significantly fewer distance computations than a sequential scan of the entire dataset.

In some applications, large-scale reference databases grow with time. Hence the new entities need to be added incrementally. For example, some applications would require the addition of queries that have no matching records in the reference database already. To facilitate those, we need to extend the Kd-tree accordingly without repeating the embedding process that creates the initial tree. However, growing a Kd-tree can be a heuristic procedure where the tree could become unbalanced. Therefore, we can explore alternative tree structures such as R-tree that are robust against dynamic data.

7 RELATED WORK

Indexing techniques are recognized as a crucial component for improving the efficiency and the scalability of the ER solutions. There exist a variety of indexing techniques which are also known as blocking or filtering. Papadakis *et al.* [26] and Christen *et al.* [8] presented surveys that include well-known indexing methods such as standard blocking, suffix array, q-gram blocking, sorted-neighbourhood, canopy clustering, and string-map based methods. We only survey a few closely related works since many of these existing methods are orthogonal to the focus of this paper.

Canopy clustering uses cheap comparison metrics to group similar records into overlapping canopies and creates blocks from records that share a common-canopy [25]. The method depends

on global threshold values, and that reduces its flexibility. It also uses similarity measures such as TF-IDF and Cosine distance that can be computationally expensive [8].

Mapping-based indexing methods map records to objects in a Euclidean space, preserving the original distances between them. Jin *et al.* [21] proposed the Stringmap algorithm that maps records into a similarity preserving Euclidean space (with dimension between 15-20). Similar pairs are determined by building an R-tree. Stringmap has linear complexity, but it requires tuning several parameters. Moreover, the performance of such approaches tends to decrease with more than 20 dimensions. In contrast, our method operates in much smaller dimensions.

The Double Embedding scheme [1] uses two-dimensions, K and K' for embedding records such that $K' < K$. Similarity joins are performed in the metric-space using a Kd-tree and nearest-neighbour search to find candidate matches. The method is faster than the Stringmap algorithm. However, it attempts to keep the embedding contractive by increasing the distance computations and is not suitable for large-scale data.

Another metric-space indexing technique utilised an M-tree to produce complete and efficient ER results. The cost of the method and the quality of the results have remained similar to existing indexing techniques [2]. However, it does not scale for large-scale data since it has a single step that combines the indexing, comparison, and classification steps.

In summary, the existing mapping-based indexing techniques have two components. The first is to map the records into a metric-space. The second is to perform similarity joins in the metric-space using a tree-like indexing structure.

These methods were developed offline by applying a mapping technique to map all the records into a multidimensional metric-space. The spatial mapping of records acts as a filtering step before the actual record matching. Hence, the focus is on matching similar records without grouping them into blocks. However, none of them handles new records by reusing the properties of an existing multidimensional metric-space. Therefore, query matching is not achievable. In this paper, similar ideas are significantly extended to accommodate the benefits of the Euclidean space for efficient query matching.

8 CONCLUSIONS

Indexing techniques reduce the pairwise comparisons in ER solutions. Many existing mapping-based indexing techniques work in offline mode with fixed-size databases. Hence, these techniques are not suitable for applications that require real-time query matching, especially if it involves big, fast, or streaming data. Our method investigated the query matching problem in ER by using spatial mapping of records into a Euclidean space. We aimed to develop an indexing approach for a fast query process within a short time, returning as many as potential matches. The proposed method proved fast running time as well as scalability along with the data size. The use of vectors in the Euclidean space that represent records allowed fast query matching with comparable accuracy. Many directions are ahead for future work. First, we plan to extend the Em-K indexing method to be parallel. Second, the current Em-K indexing can be extended to other forms of ER problems, such as querying against a dynamic database or iterative ER [18].

REFERENCES

[1] Noha Adly. 2009. Efficient record linkage using a double embedding scheme.. In *DMIN*, Vol. 48. 274–281.

[2] Özgür Akgün, Alan Dearle, Graham Kirby, and Peter Christen. 2018. *Using metric space indexing for complete and efficient record linkage.* 89–101. https://doi.org/10.1007/978-3-319-93040-4_8

[3] Sunil Arya, David M. Mount, Nathan S. Netanyahu, Ruth Silverman, and Angela Y. Wu. 1998. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *J. ACM* 45, 6 (Nov. 1998), 891–923. https://doi.org/10.1145/293347.293348

[4] Chris Chatfield and Alexander Collins. 1981. *Introduction to multivariate analysis.* CRC Press.

[5] Edgar Chávez, Gonzalo Navarro, Ricardo Baeza-Yates, and José Luis Marroquín. 2001. Searching in metric spaces. *ACM Comput. Surv.* 33, 3 (Sept. 2001), 273–321. https://doi.org/10.1145/502807.502808

[6] H. Chen, W. Chung, J.J. Xu, G. Wang, Y. Qin, and M. Chau. 2004. Crime data mining: a general framework and some examples. *Computer* 37, 4 (2004), 50–56. https://doi.org/10.1109/MC.2004.1297301

[7] Peter Christen. 2012. *Further topics and research directions.* Springer Berlin Heidelberg, Berlin, Heidelberg, 209–228. https://doi.org/10.1007/978-3-642-31164-2_9

[8] Peter Christen. 2012. A Survey of indexing techniques for scalable record linkage and deduplication. *IEEE Transactions on Knowledge and Data Engineering* 24, 9 (2012), 1537–1555.

[9] Peter Christen and Karl Goiser. 2007. *Quality and complexity measures for data linkage and deduplication.* Springer Berlin Heidelberg, Berlin, Heidelberg, 127–151. https://doi.org/10.1007/978-3-540-44918-8_6

[10] Peter Christen and Dinusha Vatsalan. 2013. Flexible and extensible generation and corruption of personal data. In *Proceedings of the 22nd ACM International Conference on Information and Knowledge Management* (San Francisco, California, USA) *(CIKM '13).* Association for Computing Machinery, New York, NY, USA, 1165–1168. https://doi.org/10.1145/2505515.2507815

[11] Uwe Draisbach and Felix Naumann. 2011. A generalization of blocking and windowing algorithms for duplicate detection. *Proceedings - 2011 International Conference on Data and Knowledge Engineering, ICDKE 2011*, 18 – 24. https://doi.org/10.1109/ICDKE.2011.6053920

[12] Christos Faloutsos and King-Ip Lin. 1995. FastMap: a fast algorithm for indexing, data-mining and visualization of traditional and multimedia datasets. In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data* (San Jose, California, USA) *(SIGMOD '95).* Association for Computing Machinery, New York, NY, USA, 163–174. https://doi.org/10.1145/223784.223812

[13] Patrick Groenen and Michel Velden. 2016. Multidimensional scaling by majorization: a review. *Journal of Statistical Software* 73 (09 2016). https://doi.org/10.18637/jss.v073.i08

[14] Samudra Herath, Matthew Roughan, and Gary Glonek. 2021. Generating name-Like vectors for testing large-scale entity resolution. *IEEE Access* 9 (2021), 145288–145300. https://doi.org/10.1109/ACCESS.2021.3122451

[15] Gisli Hjaltason and Hanan Samet. 2003. Index-driven similarity search in metric spaces. *ACM Transactions on Database Systems (TODS)* 28 (12 2003), 517–580. https://doi.org/10.1145/958942.958948

[16] Matthew A. Jaro. 1989. Advances in record-linkage methodology as applied to matching the 1985 census of Tampa, Florida. *J. Amer. Statist. Assoc.* 84, 406 (1989), 414–420. https://doi.org/10.1080/01621459.1989.10478785

[17] Pegah Kamousi, Sylvain Lazard, Anil Maheshwari, and Stefanie Wuhrer. 2016. Analysis of farthest point sampling for approximating geodesics in a graph. *Computational Geometry* 57 (2016), 1–7. https://doi.org/10.1016/j.comgeo.2016.05.005

[18] Hung-sik Kim and Dongwon Lee. 2010. HARRA: fast iterative hashed record linkage for large-scale data collections. 525–536. https://doi.org/10.1145/1739041.1739104

[19] Joseph Kruskal. 1964. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika* 29, 1 (01 Mar 1964), 1–27.

[20] Neeraj Kumar, li Zhang, and Shree Nayar. 2008. What is a good nearest neighbors algorithm for finding similar patches in images? 364–378. https://doi.org/10.1007/978-3-540-88688-4_27

[21] Chen Li, Liang Jin, and Sharad Mehrotra. 2006. Supporting efficient record linkage for large data sets using mapping techniques. *World Wide Web* 9 (12 2006), 557–584. https://doi.org/10.1007/s11280-006-0226-8

[22] Huizhi Liang, Yanzhe Wang, Peter Christen, and Ross Gayler. 2014. Noise-tolerant approximate blocking for dynamic real-time entity resolution. In *Advances in Knowledge Discovery and Data Mining*, Vincent S. Tseng, Tu Bao Ho, Zhi-Hua Zhou, Arbee L. P. Chen, and Hung-Yu Kao (Eds.). Springer International Publishing, Cham, 449–460.

[23] Mark P J Van Der Loo. 2014. The STRINGDIST package for approximate string matching. *The R Journal* 6, 1 (2014), 111–122.

[24] Arturas Mazeika and Michael Böhlen. 2006. Cleansing databases of misspelled proper nouns. https://doi.org/10.5167/uzh-56109

[25] Andrew McCallum, Kamal Nigam, and Lyle H. Ungar. 2000. 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* (Boston, Massachusetts, USA) *(KDD '00).* Association for Computing Machinery, New York, NY, USA, 169–178. https://doi.org/10.1145/347090.347123

[26] George Papadakis, Dimitrios Skoutas, Emmanouil Thanos, and Themis Palpanas. 2020. Blocking and filtering techniques for entity resolution: A survey. *ACM Comput. Surv.* 53, 2, Article 31 (March 2020), 42 pages. https://doi.org/10.1145/3377455

[27] Alieh Saeedi, Eric Peukert, and Erhard Rahm. 2017. Comparative evaluation of distributed clustering schemes for multi-source entity resolution. https://doi.org/10.1007/978-3-319-66917-5_19

[28] Vin de Silva and Joshua B. Tenenbaum. 2002. Global versus local methods in nonlinear dimensionality reduction. In *Proceedings of the 15th International Conference on Neural Information Processing Systems (NIPS'02)*. MIT Press, Cambridge, MA, USA, 721–728.

[29] Giovanni Simonini, George Papadakis, Themis Palpanas, and Sonia Bergamaschi. 2018. Schema-agnostic progressive entity resolution. *IEEE Transactions on Knowledge and Data Engineering* 31, 6 (2018), 1208–1221.

[30] Douglas Talbert and Doug Fisher. 2000. An empirical analysis of techniques for constructing and searching k-dimensional trees. 26–33. https://doi.org/10.1145/347090.347100

[31] Liying Tang and Mark Crovella. 2003. Virtual landmarks for the Internet. In *Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement* (Miami Beach, FL, USA) *(IMC '03)*. Association for Computing Machinery, New York, NY, USA, 143–152. https://doi.org/10.1145/948205.948223

[32] Gang Wang, Hsinchun Chen, and Homa Atabakhsh. 2004. Automatically detecting deceptive criminal identities. *Commun. ACM* 47, 3 (March 2004), 70–76. https://doi.org/10.1145/971617.971618