

An Advanced Entity Resolution in Data Lakes: First Steps

Lamisse F. Bouabdelli^{1,2}^a, Fatma Abdelhedi²^b, Slimane Hammoudi³^c
and Allel Hadjali¹^d

¹*LIAS Laboratory, ISAE-ENSMA, Poitiers, France*

²*CBI² Research laboratory, Trimane, Paris, France*

³*ESEO, Angers, France*

Keywords: Data Lakes, Data Quality, Entity Resolution, Entity Matching, Machine Learning.

Abstract: Entity Resolution (ER) is a critical challenge for maintaining data quality in data lakes, aiming to identify different descriptions that refer to the same real-world entity. We address here the problem of entity resolution in data lakes, where their schema-less architecture and heterogeneous data sources often lead to entity duplication, inconsistency, and ambiguity, causing serious data quality issues. Although ER has been well studied both in academic research and industry, many state-of-the-art ER solutions face significant drawbacks. Existing ER solutions typically compare two entities based on attribute similarity, without taking into account that some attributes contribute more significantly than others in distinguishing entities. In addition, traditional validation methods that rely on human experts are often error-prone, time-consuming, and costly. We propose an efficient ER approach that leverages deep learning, knowledge graphs (KG), and large language models (LLM) to automate and enhance entity disambiguation. Furthermore, the matching task incorporates attribute weights, thereby improving accuracy. By integrating LLM for automated validation, this approach significantly reduces the reliance on manual expert verification while maintaining high accuracy.

1 INTRODUCTION

The exponential growth in volume, velocity, and variety of data has introduced the concept of Big Data, which has significantly transformed how organizations store, process, and analyze information. To manage these large-scale heterogeneous datasets, organizations have adopted data lakes, scalable storage systems designed to ingest structured, semi-structured, and unstructured data in its raw format without requiring a predefined schema. This schema-less architecture offers flexibility and scalability, making data lakes attractive solutions for enterprises.


However, this type of architecture leads to numerous data quality issues due to duplicate records, inconsistencies, and variations in data representation across multiple sources. These issues impact the ac-


curacy of data analysis, leading to poor decision-making. Therefore, an efficient entity resolution approach is needed in data lake contexts.


Entity resolution (ER) (Barlaug and Gulla, 2021), is the process of determining whether two entities refer to the same real-world entity (Christophides et al., 2020) (Christen, 2012). The term *entity* refers to a distinct and identifiable unit that represents an object, a person, a place, or a concept of the real world. An entity has attributes that describe its characteristics. The term *resolution* is used because ER is fundamentally a decision-making process to resolve the question: Do the descriptions refer to the same or different entities? (Talbur, 2011). ER is also defined as "the process of identifying and merging records judged to represent the same real-world entity" (Benjelloun et al., 2007).


Current ER solutions face limitations in matching since they do not take into account the weights of the attributes, leading to mismatched possibility. In addition, they rely on manual validation, making the process expensive and time consuming.

In this paper, we describe a pipeline of an efficient entity resolution approach for data lakes. Our propo-

^a <https://orcid.org/0009-0002-9010-3128>

^b <https://orcid.org/0000-0003-2522-3596>

^c <https://orcid.org/0000-0002-9086-6793>

^d <https://orcid.org/0000-0002-4452-1647>

sition leverages deep learning, knowledge graphs for improved matching accuracy, on the one hand, and LLM for automated validation, on the other hand. In Section 2, we motivate our research work. Then, we review existing entity resolution techniques in both academia and industry in Section 3. Our entity resolution approach is explicitly discussed in Section 4. Lastly, in Section 5, we conclude by summarizing the key elements of our approach and highlighting future directions for improving entity resolution in large-scale environments.

2 CONTEXT

Today, organizations rely on their data for decision making, utilizing advanced analytics, machine learning, and business intelligence (BI) tools to gain strategic insights and operational efficiency.

Figure 2 illustrates an architecture case of an organization where data originate from multiple sources. These data are stored in a data lake, then these data are destined to be cleaned in order to be ingested into a data warehouse with aim of using it for different analytics purposes for the intention of making optimal decisions.

However, since these data come from heterogeneous sources, they can vary in structure, format, and semantics. These variations lead to entity inconsistencies, duplicate records, missing attributes, and lack of standardization, all of which degrade the accuracy and reliability of analytical outputs, leading to incorrect decision making. Therefore, entity resolution is mandatory.

The challenge in data quality management is entity ambiguity, which occurs when multiple representations of the same real-world entity exist within or across datasets. Figure 1(a) shows an example of two records from two different sources that refer to the same person. Figure 1 (b) represents an example of two records from two different sources that are not the same person, even though they have quite similar names, same addresses, and dates of birth that differ only by two transposed numbers in the year. Assuming these are the same person would be a false positive. For this reason, entity resolution is crucial and must take into consideration all possible cases to correctly match entities, because failing to resolve these ambiguities can lead to erroneous insights and operational inefficiencies.

Since organizations rely on data for important decision-making, the need for a robust entity resolution solution has never been more critical. The entity resolution process, which consists of identify-

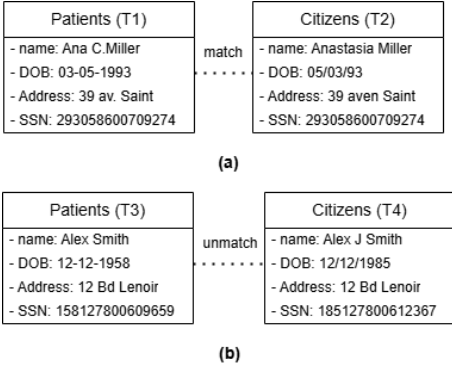


Figure 1: Example of an entity resolution problem.

ing, matching, and merging records that refer to the same real-world entity, is essential for maintaining data integrity, consistency, and reliability. In high-stakes domains such as healthcare, finance, and e-commerce, errors in entity resolution can have severe consequences, from incorrect patient records leading to misdiagnoses, to fraudulent financial transactions, or misattribute customer data affecting business decisions. Hence, the main question is: How can we improve entity resolution in schema-less data lakes?

To address these challenges, our research focuses on improving data quality in data lakes through an advanced entity resolution proposition. Our proposition uses advanced techniques, including deep learning models and knowledge graphs, to enhance similarity detection and capture relationships, and LLM, to automate the validation process.

By integrating these techniques, we aim to propose a robust, scalable, and automated entity resolution approach capable of handling large-scale, heterogeneous datasets while improving overall data quality. This work is particularly relevant for organizations looking to optimize data lake usability and improve business intelligence insights for optimal decision-making.

3 RELATED WORK

Entity resolution has been a key area of interest both in academic research (Barlaug and Gulla, 2021) (Christen, 2012) and industry, evolving significantly from traditional similarity measures to machine learning techniques that have shown an improvement in matching performance (Christophides et al., 2020). By the late 2010s, deep learning became a key area of research in data matching (Peeters et al., 2024) (Mudgal et al., 2018) (Li et al., 2020). Other research has studied ER using graph-based methods (Yao et al., 2022), and more recently experimented with LLM (Peeters




Figure 2: An entity resolution architecture.

et al., 2024).

The industry has proposed many entity resolution solutions using machine learning and artificial intelligence. Among these solutions, Sensing (Sensing, 2025), designed for entity matching, combines ML clustering and AI. Quantexa (Quantexa, 2025) employs ML and AI techniques, offers entity linkage; however, its reliance on complex graph structures poses implementation challenges for non-expert users. AWS Glue (AWS, 2017), a cloud-based ER solution, integrates entity resolution within broader ETL workflows. Its scalability and seamless integration with AWS services make it a powerful tool. DataWalk (Datawalk, 2025), on the other hand, is a unified graph and AI platform for data management, analysis, and investigative intelligence, which includes entity resolution software.

Despite the advancements offered by these tools, industry ER solutions still face key limitations. A critical drawback is the lack of attribute weighting, where all entity attributes are treated equally despite varying levels of significance, which can lead to suboptimal matching results. Furthermore, the validation phase often relies on manual intervention, thereby increasing operational costs and time. This human dependence not only reduces the efficiency of these solutions, but also introduces human error. Due to these issues, there is a clear need for improvement in industrial ER tools to better address these challenges. Enhanced attribute weighting mechanisms and the use of reliable automated validation could significantly refine the accuracy and efficiency of entity resolution in industry.

Entity resolution has been a main focus in research for decades and is still receiving attention (Christophides et al., 2020). It started with domain experts matching entities by hand (Fellegi and Sunter, 1969). Now, with advances in technology, machine learning-based approaches have been introduced, using supervised and unsupervised learning techniques to improve ER (Christophides et al., 2020). Methods

such as Support Vector Machines (Cortes and Vapnik, 1995), (Bilenko and Mooney, 2003) classify entity pairs based on engineered similarity features, while Random Forests (Breiman, 2001) employ ensemble learning to improve classification performance. However, these models require extensive feature engineering and struggle with unseen entity variations, limiting their adaptability to large and evolving datasets. Transformer and pre-trained models like BERT (Devlin et al., 2018) and RoBERTa (Liu et al., 2019) revolutionized natural language processing. Studies have explored entity matching using pre-trained models (Paganelli et al., 2024) (Li et al., 2021). More recently, deep learning models have significantly advanced entity resolution by capturing contextual dependencies between entity attributes. DeepMatcher (Mudgal et al., 2018) applies bidirectional LSTM (Hochreiter and Schmidhuber, 1997) with attention mechanisms to learn entity similarity from labeled data, while Ditto (Li et al., 2020) uses transformer-based architectures to fine-tune pre-trained models on ER tasks. Ditto brings some optimizations that require domain knowledge. These deep learning-based methods are based on text sequences for matching. They use different methods for attribute embedding and attribute similarity representation. Furthermore, HierGAT (Hierarchical Graph Attention Networks) (Yao et al., 2022) enhances entity matching by incorporating graph-based relationships, demonstrating the potential of graph neural networks (GNNs) for ER problems. Despite their improvements in precision and recall, deep learning-based methods often overlook the importance of attribute weighting and struggle with explainability, posing challenges for real-world adoption. The advent of LLM such as Llama and GPT has further pushed the boundaries of ER by enabling zero-shot and few-shot learning for entity matching tasks (Peeters et al., 2024). Although LLM have shown strong performance, their effectiveness remains highly dependent on domain-specific fine-tuning and prompt engineering, making them compu-

tationally expensive and less adaptable to structured relational datasets. Moreover, existing LLM-based approaches do not inherently model inter-entity relationships, which limits their applicability in graph-based ER scenarios.

To clarify, an LLM that performs matching based solely on textual attributes might miss the underlying relationships between entities. For instance, consider a father and son who share the same last name and home address. An LLM could mistakenly classify them as the same person due to the high textual similarity of their attributes. However, the crucial relationship (father-son) indicates they are related but distinct individuals. This relational nuance cannot be captured by the LLM alone. In contrast, a knowledge graph can explicitly represent such relationships, enabling the system to recognize them as separate entities. This example demonstrates why relying exclusively on LLM can be problematic in graph-based ER settings: LLM lack explicit, structured mechanisms to represent and reason over inter-entity relationships.

In contrast, research efforts have also explored rule-based methods (Singh et al., 2017) that require designing rules and setting thresholds and crowdsourcing-based ER methods (Wang et al., 2012), which require extensive manual intervention or rely on human annotators to validate entity matches. As ER continues to evolve, our research focus on hybrid approaches that combine deep learning, knowledge graph, and pre-trained LLM, leveraging the strengths of each paradigm to improve entity resolution across diverse real-world datasets.

4 METHODOLOGY AND TECHNIQUES

This section introduces our proposed entity resolution approach. Figure 3 illustrates the pipeline of our proposed method. This pipeline is supposed to be placed within the architectural framework shown in Figure 2, especially between the data lake and the data warehouse.

Our approach is inspired by established techniques in the literature (Christen, 2012), but introduces key adaptations to improve the entity resolution process. The process consists of four main steps: 1) Pre-processing ensures data quality by standardizing formats, correcting mistyping errors, handling missing values, and normalizing variations. 2) Blocking aims to reduce computational complexity by grouping similar records in the same block to limit entity comparisons to subsets. 3) Matching compares records within the same block in order to iden-

tify records that correspond to the same real-world entity. 4) Validation is traditionally carried out by domain experts, which is often time consuming and costly. To address this, we propose an automated validation mechanism that significantly reduces manual effort.

The following subsections provide a detailed explanation of each phase of the pipeline, including the specific techniques and methodology used.

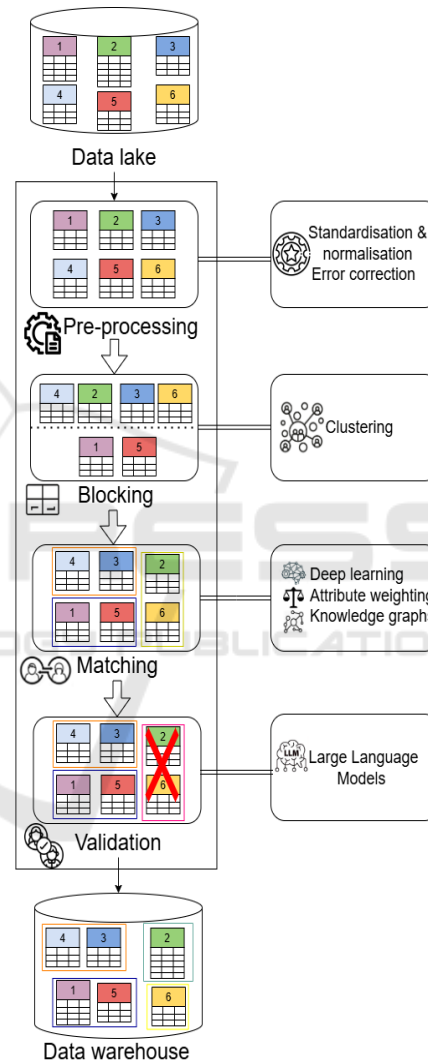


Figure 3: Pipeline of the proposed approach.

4.1 Pre-Processing

Pre-processing is a critical step because it ensures quality of data, which is essential for improving entity resolution. It requires several key operations including standardization, where data formats such as dates and addresses are unified, correcting mistyping

errors, and identifying missing values. Additionally, linguistic normalization is applied to unify abbreviations, acronyms, and variations of entity names, plus special character removal and the elimination of unnecessary punctuations, symbols, and whitespace.

Given that real-world data are often noisy and incomplete, we aim to improve the data quality by assuring accuracy (ensuring that data correctly reflect real-world entities), consistency (ensuring that data are harmonized and uniformed across multiple sources), correctness (verify data validity), and completeness (assessing whether all the essential information is present). Completeness is further categorized into: total completeness means no missing data, partial completeness some missing data, but it will not affect the processes and the information remains exploitable, critical completeness where essential data are missing.

The goal of pre-processing is to enhance data quality for the next steps. The output of the pre-processing is clean data, for the purpose of reducing the number of sets for the matching, for this reason we introduce our next phase.

4.2 Blocking

Blocking is an optimization step designed to reduce the number of comparisons between entity pairs, thus significantly reducing computational costs. Instead of evaluating all possible entity pairs, blocking groups similar records together in the same block. This ensures that only the most relevant subsets are considered for a detailed matching (Christophides et al., 2020).

Various blocking techniques have been explored in the literature (Christophides et al., 2020)(Skoutas et al., 2019)(Paganelli et al., 2024), each with its own advantages. In our approach, we are going to use unsupervised clustering as an effective method to group similar records based on their attributes.

After grouping potentially similar records in blocks and reducing computational complexity as a result, limiting entity comparisons to subsets that are ready for the next step.

4.3 Matching

The matching phase constitutes the most critical step in entity resolution, as it seeks to identify records that correspond to the same real-world entity, despite variations in their descriptions, a phenomenon known as synonymy, as illustrated in Figure 1(a). In contrast, it is equally crucial to differentiate records that may exhibit similar attributes but actually represent distinct

entities, a challenge called homonymy or entity collision, as shown in Figure 1(b).

Our proposed matching approach incorporates attribute weighting, recognizing that some attributes contribute more significantly than others in distinguishing entities. We acknowledge that previous ER approaches have incorporated attribute importance through weighted similarity. Notably, recent graph-based models like (Yao et al., 2022) employ attention mechanisms to identify the most discriminative attributes, our approach builds upon these by introducing an explicit, tunable weighting mechanism. This mechanism allows for greater control and transparency compared to deep models where attribute importance is learned implicitly. To further enhance matching accuracy, our approach combines deep learning techniques for measuring similarity between attribute values with knowledge graphs that capture the relationships between entities that are likely to match. This hybrid approach will ultimately improve ER by ensuring a more context-aware, semantically enriched, and structurally informed matching process, leading to higher precision and reduced false positives.




Figure 4: Example of entity resolution.

4.3.1 Problem Formulation

Figure 4 illustrates a scenario in which existing entity resolution solutions may incorrectly merge two distinct entities due to high similarity in certain attributes. Specifically, Entity 1 (T3) in the patient table and Entity 2 (T4 in the citizen table and T5 in the employee table) share common data points, such as address and name variations, making them appear as potential duplicates. The question arises: *How does our proposed matching approach differ and why is it more effective?*

To formalize the problem, we define the following.

- A dataset consisting of multiple tables T , where each table contains a set of attributes denoted as:

$$A = \{a_1, a_2, \dots, a_n\}. \quad (1)$$

- Attributes, such as in our example name, date of birth, address, and Social Security Number (SSN).

- A weight function $w : A \rightarrow [0, 1]$ that assigns a weight to each attribute based on its discriminative power. For instance, attributes like SSN have a high weight due to their uniqueness:

$$\begin{aligned} w(\text{SSN}) &= 0.9, \\ w(\text{name}) &= 0.7, \\ w(\text{DOB}) &= 0.5, \\ w(\text{address}) &= 0.4. \end{aligned} \quad (2)$$

Matching Computation: Given two records T_i and T_j , we calculate a similarity score for each attribute using deep learning techniques to measure the degree of correspondence between attribute values. This results in a similarity vector:

$$\text{match}(T_3, T_4) = \{s_1, s_2, \dots, s_n\}. \quad (3)$$

For example:

$$\text{match}(T_3, T_4) = \{0.6, 0.8, 1.0, 0.4\}. \quad (4)$$

where each s_k represents the similarity score for the attribute a_k .

To compute the final matching score, we apply a weighted sum:

$$S(T_3, T_4) = \frac{1}{n} \sum_{k=1}^n w(a_k) \cdot s_k. \quad (5)$$

Alternatively, instead of a simple summation, we propose using the Skyline operator (Borzsony et al., 2001), which is considered an optimization solution that selects non-dominated matches based on Pareto optimality.

By incorporating attribute weighting and deep learning-based similarity computation. We aim that our approach will significantly reduce false positives while improving the precision of entity resolution.

4.3.2 Capturing Relationships with Knowledge Graphs

Figure 4 illustrates an example in which existing entity resolution methods struggle, often erroneously matching two entities that are, in fact, distinct. However, our method goes beyond similarity matching by using knowledge graphs to detect relationships between entities rather than incorrectly merging them.

By integrating knowledge graphs, our approach captures semantic relationships between entities. In this example, instead of falsely concluding that Entity 1 and Entity 2 are the same person, we find a related relationships, they are likely father and son. Social Security Numbers (SSN) differ, but name, address, and other attributes share similarities, which can mislead conventional entity resolution methods.

To model this, we define $E = \{e_1, e_2, \dots, e_n\}$ as the set of entities. R as the set of possible relationships between entities, where each relationship is defined as a directed edge $r(e_i, e_j)$ in the knowledge graph. A similarity function $S(T_i, T_j)$ that computes the weighted similarity between records, capturing both direct attribute matches and inferred relationships.

Using this structure, our method assigns relationship probabilities instead of merely merging entities. The system recognizes that while Entity 1 and Entity 2 are distinct, they are related, thus preventing false positives in entity resolution.

We believe that by combining deep learning techniques for attribute matching with knowledge graphs for relationship inference, our approach will achieve higher accuracy in distinguishing similar but distinct entities while assuring the preservation of important relationships rather than erroneous resolved entities. Plus scalability in handling complexities in real-world data.

Lastly, after finding entities that match and for the purpose of affirming if the resolved entities are correctly matched, we present our next step for entity validation.

4.4 Validation

The validation phase aims to verify that the resolved entities match correctly. Traditionally, validation relies on domain experts to manually verify. This approach, while reliable, is highly time consuming, costly, and prone to human errors, especially when dealing with large-scale datasets.

To overcome these limitations, we propose an automated validation mechanism using LLM to validate the resolved entities without human intervention. Our approach utilizes the reasoning and contextual understanding capabilities of LLM to assess whether two records represent the same entity.

We are aware that LLM can be used during the matching phase. However, we deliberately restrict their use to the validation phase because of considerations of cost, scalability, and explainability. Running an LLM on every candidate pair during matching would be computationally expensive and inefficient, especially when processing millions of comparisons. In contrast, using LLM only on a reduced subset of record pairs, those that survived earlier blocking and matching a better balance between accuracy and performance. This approach allows us to benefit from LLM sophisticated reasoning capabilities. LLM validation step provides a final layer of confidence by validating only the top-ranked candidate pairs.

By automating validation, we eliminate the reliance on domain experts for this task, consequently reducing human effort and operational costs. The approach is highly scalable, capable of efficiently validating millions of records, which makes it well suited for large-scale entity resolution. Using the contextual reasoning capabilities of LLM, our method aims to ensure high accuracy by minimizing false positives and false negatives. Furthermore, the execution speed of our validation mechanism is faster than manual methods, enabling real-time or near-real-time verification. IN summary, by integrating deep learning, knowledge graphs, and LLM, our entity resolution approach aims to ensure a more efficient, scalable, and reliable validation process.

5 CONCLUSION AND FUTURE WORK

Data quality is a critical challenge in data lakes. Therefore, entity resolution is crucial to enhance data quality which is essential for making optimal decisions.

In this paper, we propose a novel entity resolution approach designed to improve data quality, scalability, and automation in data lakes. Our solution uses deep learning, to improve entity matching, knowledge graphs, to capture relationships between entities and LLM to reduce human intervention in the validation phase.

Our approach presents a potentially effective improvement to existing entity resolution solutions, but its true performance and efficiency can only be validated through real-world implementation and experimentation.

Since our work is currently a theoretical proposition, our next step is to implement this approach and conduct a comprehensive evaluation against existing solutions. We aim to demonstrate its effectiveness in the real-world and ultimately contribute to the advancement of entity resolution.

While our approach focuses on the identification of duplicate entities, we acknowledge that the subsequent step data fusion (merging duplicate records into unified representations) is not addressed in this paper. Data fusion is a critical and non trivial component of the ER pipeline, and we plan to investigate scalable and context-aware fusion strategies as part of future work.

However, we note that data fusion has already been explored in previous research efforts (Abdelhedi et al., 2022a) (Abdelhedi et al., 2022b) (Abdelhedi et al., 2021), where our team explored merging dupli-

cate records in data lakes using ontology-driven integration. Building upon such foundations, our future efforts will aim to incorporate a robust, semantically informed fusion module to complete the ER pipeline.

ACKNOWLEDGEMENTS

The authors acknowledge Professor Gilles Zurfluh for his invaluable advice, insightful ideas, and time for this work. His expertise and thoughtful advice have been critical in shaping the direction of this work.

REFERENCES

- Abdelhedi, F., Jemmali, R., and Zurfluh, G. (2021). Ingestion of a data lake into a nosql data warehouse: The case of relational databases. In *KMIS*, pages 64–72.
- Abdelhedi, F., Jemmali, R., and Zurfluh, G. (2022a). Data Ingestion from a Data Lake: The Case of Document-oriented NoSQL Databases. In Filipe, J., Smialek, M., Brodsky, A., and Hammoudi, S., editors, *Proceedings of the 24th International Conference on Enterprise Information Systems - ICEIS 2022* ; ISBN 978-989-758-569-2 ; ISSN 2184-4992, volume 1: ICEIS, pages 226–233, Online Streaming, France. SCITEPRESS : Science and Technology Publications.
- Abdelhedi, F., Jemmali, R., and Zurfluh, G. (2022b). DLToDW: Transferring Relational and NoSQL Databases from a Data Lake. *SN Computer Science*, 3(5):article 381.
- AWS (2017). Aws glue. <https://aws.amazon.com/fr/glue/>.
- Barlaug, N. and Gulla, J. A. (2021). Neural networks for entity matching: A survey. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 15(3):1–37.
- Benjelloun, O., Garcia-Molina, H., Gong, H., Kawai, H., Larson, T. E., Menestrina, D., and Thavisomboon, S. (2007). D-swoosh: A family of algorithms for generic, distributed entity resolution. In *27th International Conference on Distributed Computing Systems (ICDCS'07)*, pages 37–37. IEEE.
- Bilenko, M. and Mooney, R. (2003). Adaptive duplicate detection using learnable string similarity measures. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 39–48.
- Borzsony, S., Kossmann, D., and Stocker, K. (2001). The skyline operator. In *Proceedings 17th international conference on data engineering*, pages 421 – 430. IEEE.
- Breiman, L. (2001). Random forests. *Machine learning*, 45:5–32.
- Christen, P. (2012). *Data Matching*. Springer: Data-centric systems and applications.

- Christophides, V., Efthymiou, V., Palpanas, T., Papadakis, G., and Stefanidis, K. (2020). An overview of end-to-end entity resolution for big data. *ACM Computing Surveys (CSUR)*, 53(6):1–42.
- Cortes, C. and Vapnik, V. (1995). Support-vector networks. *Machine learning*, 20:273–297.
- Datawalk (2025). Data walk entity resolution. <https://datawalk.com/solutions/entity-resolution/>.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, page 15.
- Fellegi, I. P. and Sunter, A. B. (1969). A theory for record linkage. *Journal of the American statistical association*, 64(328):1183–1210.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Li, Y., Li, J., Suhara, Y., Doan, A., and Tan, W.-C. (2020). Deep entity matching with pre-trained language models. *Proceedings of the VLDB Endowment*, 14:50–60.
- Li, Y., Li, J., Suhara, Y., Wang, J., Hirota, W., and Tan, W.-c. (2021). Deep entity matching: Challenges and opportunities. *Journal of Data and Information Quality (JDIQ)*, 13:1–17.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). Roberta: A robustly optimized bert pre-training approach. *arXiv preprint arXiv:1907.11692*.
- Mudgal, S., Li, H., Rekatsinas, T., Doan, A., Park, Y., Krishnan, G., Deep, R., Arcaute, E., and Raghavendra, V. (2018). Deep learning for entity matching: A design space exploration. In *Proceedings of the 2018 international conference on management of data*, pages 19–34.
- Paganelli, M., Tiano, D., and Guerra, F. (2024). A multi-facet analysis of bert-based entity matching models. *The VLDB Journal*, 33(4):1039–1064.
- Peeters, R., Steiner, A., and Bizer, C. (2024). Entity matching using large language models. *arXiv preprint arXiv:2310.11244*.
- Quantexa (2025). Quantexa. <https://www.quantexa.com/fr/>.
- Senzing, I. (2025). Senzing – entity resolution software. <https://senzing.com/>.
- Singh, R., Meduri, V., Elmagarmid, A., Madden, S., Papotti, P., Quiané-Ruiz, J.-A., Solar-Lezama, A., and Tang, N. (2017). Generating concise entity matching rules. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 1635–1638.
- Skoutas, D., Thanos, E., and Palpanas, T. (2019). A survey of blocking and filtering techniques for entity resolution. *arXiv preprint arXiv:1905.06167*.
- Talbur, J. (2011). *Entity Resolution and Information Quality*. Elsevier.
- Wang, J., Kraska, T., Franklin, M. J., and Feng, J. (2012). Crowder: Crowdsourcing entity resolution. *arXiv preprint arXiv:1208.1927*.
- Yao, D., Gu, Y., Cong, G., Jin, H., and Lv, X. (2022). Entity resolution with hierarchical graph attention networks.

In *Proceedings of the 2022 International Conference on Management of Data*, pages 429–442.