Responsible Entity Resolution over Streaming Data

Kostas Stefanidis
Data Science Research Centre
Tampere University
Tampere, Finland
konstantinos.stefanidis@tuni.fi

Vasilis Efthymiou
Department of Informatics and Telematics
Harokopio University of Athens
Athens, Greece
vefthym@hua.gr

Tiago Brasileiro Araújo
Data Science Research Centre
Tampere University
Tampere, Finland
tiago.brasileiroaraujo@tuni.fi

I. Abstract

A. Background & Methods

Entity Resolution (ER) is a fundamental data management task that identifies and potentially merges records referring to the same real-world entity across heterogeneous sources [1], [2]. In dynamic settings, such as social media, healthcare, or finance, data arrive continuously and require timely, fair, and interpretable integration [3], [4]. While recent ER approaches achieve remarkable accuracy, they often behave as black boxes and may propagate or amplify data biases [5]. Such biases can lead to systematic disadvantages for specific groups, undermining both user trust and the reliability of downstream applications [6]. Existing fairness-aware ER approaches mainly focus on static, batch-oriented settings, while explainability is typically addressed in offline or post-hoc analyses, leaving real-time accountability and adaptive decision making largely unexplored in continuous data flows [7].

To address these challenges, we introduce X-TREATS, an explainable and fairness-aware ER framework for streaming environments. Built upon the TREATS workflow [8], X-TREATS integrates local explanation mechanisms directly into incremental, fairness-constrained entity matching. The framework employs pre-trained deep learning models for efficient comparison over micro-batches of streaming data and introduces a dual scoring mechanism to quantify both the actual influence of each attribute and its ethically desirable influence on matching decisions, defined as the expected attribute contribution under fairness constraints derived from group parity objectives. Explanations are generated using model-agnostic feature attribution methods, such as perturbation-based and SHAP-like analyses [9], and are continuously updated as new data arrive to support explainability monitoring and adaptive adjustment in the streaming ER process. Together with fairness metrics, these explanations are computed on the fly, enabling ongoing auditing, transparency, and responsible adaptation of the resolution process.

Therefore, a key contribution of X-TREATS lies in establishing an explicit feedback loop between explainability and fairness within the streaming ER pipeline. In contrast to post-hoc approaches, explanations are not treated as passive artifacts but as operational signals that actively influence matching decisions. For each candidate pair, local explanations quantify the actual influence of attributes, which are aggregated into an explanation score. This score is integrated into the fairness-aware ranking process, allowing the system to favor matches supported by interpretable and ethically consistent evidence. As new data arrive, explanation signals are continuously updated and re-incorporated into the ranking and aggregation steps, enabling ongoing auditing and adaptive adjustment of fairness constraints. Through this mechanism, X-TREATS transforms explainability into a dynamic control signal, coupling transparency and fairness to support responsible decision making under streaming constraints.

B. Results & Conclusions

Experimental evaluations on real-world streaming datasets demonstrate that our approach achieves comparable accuracy to state-of-the-art ER systems while substantially improving fairness across multiple protected groups. Using group-level measures such as True Positive Rate Parity, Positive Predictive Value Parity, and Bias@k, X-TREATS consistently reduces disparities without degrading effectiveness. Across all datasets, X-TREATS achieved Precision$\geq$0.90, while reducing Bias by up to 70%. TPRP and PPVP disparities quickly converged toward zero (the ideal fairness condition). Although transient fluctuations may occur in early increments or under skewed data distributions, TPRP and PPVP disparities quickly converge toward zero as the stream evolves, reflecting stable fairness behavior over time.

The Explanation Score, which quantifies how well the model's rationale aligns with explainability principles, increased by up to 20%, showing that explanations can serve as actionable tools for ethical monitoring and continuous model refinement. These gains indicate that explanation signals not only enhance transparency but also contribute to stabilizing the trade-off between accuracy and fairness under streaming constraints, supporting adaptive and responsible ER.

The results confirm that integrating explanation into fairness-aware ER transforms transparency from an auxiliary feature into a central mechanism for trust and equity. X-TREATS thus pioneers responsible entity resolution under streaming constraints, bridging the gap between fairness, interpretability, and real-time decision making. The key takeaway is that explainability can guide fairer matching behavior while preserving performance, paving the way toward trustworthy, human-aligned, and socially aware data integration systems.

REFERENCES

[1] V. Christophides, V. Efthymiou, and K. Stefanidis, *Entity Resolution in the Web of Data*, ser. Synthesis Lectures on the Semantic Web: Theory and Technology. Morgan & Claypool Publishers, 2015.

[2] V. Efthymiou, K. Stefanidis, and V. Christophides, "Benchmarking blocking algorithms for web entities," *IEEE Trans. Big Data*, vol. 6, no. 2, pp. 382–395, 2020.

[3] N. Shahbazi, N. Danevski, F. Nargesian, A. Asudeh, and D. Srivastava, "Through the fairness lens: Experimental analysis and evaluation of entity matching," *Proceedings of the VLDB Endowment*, vol. 16, no. 11, pp. 3279–3292, 2023.

[4] N. Fanourakis, V. Efthymiou, V. Christophides, D. Kotzinos, E. Pitoura, and K. Stefanidis, "Structural bias in knowledge graphs for the entity alignment task," in *The Semantic Web - 20th International Conference, ESWC 2023, Hersonissos, Crete, Greece, May 28 - June 1, 2023, Proceedings*, ser. Lecture Notes in Computer Science, vol. 13870. Springer, 2023, pp. 72–90.

[5] C. Fragkathoulas, V. Papanikou, D. P. Karidi, and E. Pitoura, "On explaining unfairness: An overview," in *2024 IEEE 40th International Conference on Data Engineering Workshops (ICDEW)*. IEEE, 2024, pp. 226–236.

[6] N. Shahbazi, J. Wang, Z. Miao, and N. Bhutani, "Fairness-aware data preparation for entity matching," in *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 2024, pp. 3476–3489.

[7] T. Brasileiro Araújo, V. Efthymiou, and K. Stefanidis, "Fairness and explanations in entity resolution: An overview," *IEEE Access*, vol. 13, pp. 145 127–145 143, 2025.

[8] T. B. Araújo, V. Efthymiou, V. Christophides, E. Pitoura, and K. Stefanidis, "TREATS: fairness-aware entity resolution over streaming data," *Inf. Syst.*, vol. 129, p. 102506, 2025. [Online]. Available: https://doi.org/10.1016/j.is.2024.102506

[9] N. Barlaug, "Lemon: explainable entity matching," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 8, pp. 8171–8184, 2022.