Neuro-Relational Programs: Unifying Queries and Neural Computation over Structured Data

Arie Soeteman (speaker) [1], Balder ten Cate [1], Maurice Funk [2], Benny Kimelfeld [3], Carsten Lutz [2] and Moritz Schönherr [2]

[1] ILLC, University of Amsterdam, Netherlands
{a.w.soeteman,b.d.tencate}@uva.nl
[2] Leipzig University, Germany; Center for Scalable Data Analytics and Artificial Intelligence, Germany
[3] Technion, Israel; RelationalAI, United States

We present results from [5], to appear at PODS 2027, and relate them to other approaches to machine learning on structured data. The conventional approach to deep learning over relational databases applies neural models, such as Graph Neural Networks (GNNs), to a graph representation of the database. Recent approaches instead operate on databases directly, associating tuples with embeddings and extending query mechanisms to jointly process embeddings and relational content. Inspired by these developments, we introduce Neuro-Relational Programs (NRPs), a declarative query language for relational databases whose facts carry numeric vector embeddings. NRPs extend Datalog-style rules with operations that combine, aggregate, and transform embeddings, thereby interleaving relational reasoning and learnable neural components within a single formalism. This yields a general approach to neural computation over relational data: an NRP can be read both as a query with trainable components and as a neural architecture with relational structure built in. Natural syntactic fragments of NRPs recover existing architectures and query formalisms. Zero-ary NRPs correspond to non-adaptive query algorithms [2, 6]; monadic NRPs generalize GNN-style message-passing and precisely capture Deep Homomorphism Networks [3, 4], a connection that we extend to frontier-guarded NRPs over databases with row-ids. We characterize the expressive power of unrestricted NRPs with ReLU-FFN transformations by FOCQ, an extension of first-order logic with counting interpreted over real-weighted structures, yielding a precise connection with uniform TC^0 over ordered databases. In particular, this resolves a recent question posed in the context of the more general Lara framework [1]. Together, these results establish NRPs as a broad declarative framework for querying and neural computation over relational data.

References

[1] Pablo Barceló, Nelson Higuera, Jorge Pérez, and Bernardo Subercaseaux. On the expressiveness of lara: A proposal for unifying linear and relational algebra. *Theor. Comput. Sci.*, 935:105–127, 2022.

[2] Yijia Chen, Jörg Flum, Mingjun Liu, and Zhiyang Xun. On algorithms based on finitely many homomorphism counts. *Information and Computation*, 306:105326, 2025.

[3] Takanori Maehara and Hoang NT. Deep homomorphism networks. *Advances in Neural Information Processing Systems*, 37:56076–56107, 2024.

[4] Moritz Schönherr, Balder ten Cate, Maurice Funk, Benny Kimelfeld, Carsten Lutz, and Arie Soeteman. Expressive power of deep homomorphism networks over relational databases, 2026.

[5] Arie Soeteman, Balder ten Cate, Maurice Funk, Benny Kimelfeld, Carsten Lutz, and Moritz Schönherr. Neuro-relational programs: Unifying queries and neural computation over structured data. *arXiv preprint arXiv:2606.11946*, 2026.

[6] Wei-Lin Wu. *A Study of the Expressive Power of Homomorphism Counts.* PhD thesis, University of California, Santa Cruz, USA, 2023.