Trigger Graphs and Probabilistic Equivalence: Towards Scalable and Efficient Inference and Neurosymbolic Learning

Efthymia Tsamoura[1,*,†]

[1]*Huawei Labs, Cambridge, UK*

Abstract

This keynote introduces trigger graphs, a scalable symbolic reasoning technique enabling efficient (probabilistic) Datalog reasoning over large-scale graph stores. It also presents the new equivalence semantics for probabilistic logic programming that leads to improvements of up to 42% in neurosymbolic learning.

Keywords

Symbolic reasoning, Datalog, Neurosymbolic AI,

1. Talk Summary

It has been a common belief that symbolic reasoning does not scale. However, is this still true? In this talk, I will present trigger graphs, a symbolic reasoning technique that supports exact Datalog reasoning in the order of seconds over graph stores with billions of edges. Unlike the majority of commercial and open source reasoning engines, trigger graphs avoid redundant computation during reasoning by organizing the computation in a graph-like structure. The latter allows trigger graphs to support probabilistic reasoning that is more efficient even than approximate techniques.

This year, trigger graphs became the driving force behind a new probabilistic logic programming semantics, the *equivalence semantics*. In the equivalence semantics, a probabilistic logic program induces a probability distribution over all possible equivalence relations between symbols, instead of a probability distribution over all possible subsets of probabilistic facts, as is standard in the relevant literature. We show that equivalence semantics overcomes the limitations in learning and inference of state-of-the-art neurosymbolic techniques for link prediction, rule mining, and symbolic grounding by up to 42%.

2. Biography

Efthymia (Efi) Tsamoura is a Technical Expert at Huawei Labs. From 2019 to 2025, she was a Senior Researcher at Samsung AI, Cambridge, UK. In 2016, she was awarded a prestigious early-career fellowship from the Alan Turing Institute, UK, for her work on logic and databases, and before that, she was a Postdoctoral Researcher in the Department of Computer Science of the University of Oxford. Her main research interests lie in the areas of logic, knowledge representation and reasoning, and neurosymbolic learning, while her recent outcomes involve scaling symbolic reasoning to billions of triples, as well as addressing open problems in neurosymbolic learning. Her research has been published in top-tier AI and database venues (NeurIPS, ICML, SIGMOD, VLDB, PODS, AAAI, IJCAI, etc.). In 2024, Efi was invited by the Royal Society, UK, to the Frontiers of Science on AI meeting to discuss the risks of AI and ways to address them.

Published in the Proceedings of the Workshops of the EDBT/ICDT 2026 Joint Conference (March 24-27, 2026), Tampere, Finland

efthymia.tsamoura@gmail.com (E. Tsamoura)

https://tsamoura.github.io/ (E. Tsamoura)

0009-0008-8302-1902 (E. Tsamoura)

Declaration on Generative AI

The author has not employed any Generative AI tools.