Samsung Research

Highly-Efficient Reasoning via Trigger Graphs

Efi Tsamoura

SAIC-Cambridge

| Dec 6, 2022 | SAIC-Cambridge

Disclaimer

This talk is **not** about stream reasoning…

…however, it presents a reasoning technique that can be naturally extended to reasoning in an incremental fashion.

The reasoning technique is based on a notion called **Trigger Graphs (TGs)**.

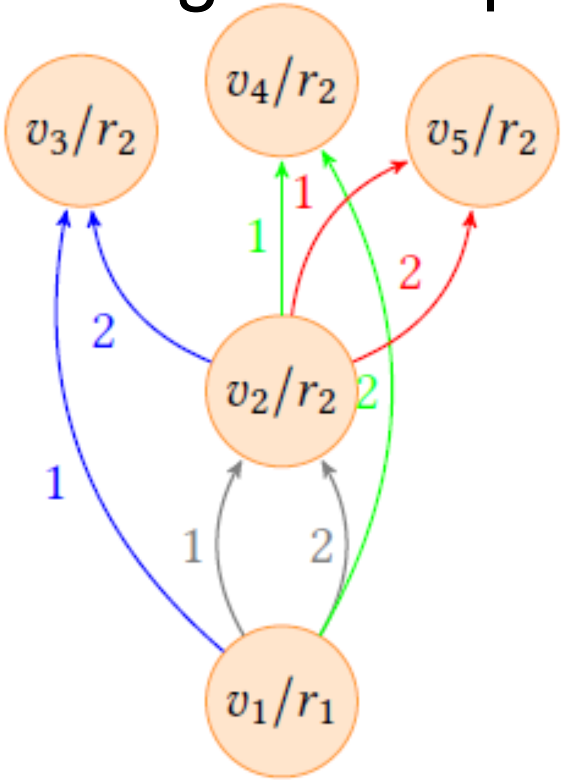

$$p(X,Y) \leftarrow e(X,Y) \qquad (r_1)$$
$$p(X,Y) \leftarrow p(X,Z) \wedge p(Z,Y) \qquad (r_2)$$

"Materializing Knowledge Bases via Trigger Graphs", **VLDB 2021**

"Probabilistic Reasoning at Scale: Trigger Graphs to the Rescue", **SIGMOD 2023**

Performance benefits of Trigger Graphs: non-probabilistic reasoning

Table 4: Datalog scenarios. Runtime is in sec and memory in MB. $*$ denotes timeout after 1h.

	VLog		RDFox		COM		GLog Runtime			GLog Memory		
Scenario	Runtime	Memory	Runtime	Memory	Runtime	Memory	No opt	m	m+r	No opt	m	m+r
LUBM-L	1.5	324	23	2301	20.4	4479	2.4	2.2	1.0	446	424	264
LUBM-LE	170.5	2725	116.6	3140	115.9	3610	17.3	17.2	16.1	1340	1310	1338
UOBM-L	7.3	1021	10	784	10	4215	2.6	2.4	2.6	335	335	342
DBpedia-L	41.6	827	64.4	3290	198.4	3878	20	19	19	1341	1352	1339
Claros-L	431	3170	2512	5491	2373.0	6453	122	118.3	119	6076	6077	6078
Claros-LE	2771.8	11 895	*	*	*	*	1040.8	1012.2	1053.9	48 464	48 474	48 455

Performance benefits of Trigger Graphs: probabilistic reasoning

Samsung Research

Table 3: Total time (default is ms) to answer the queries in LUBM010 and LUBM100 with ProbLog2 (P), Scallop (S), vProbLog (vP) and LTGs (L). Probabilities are computed via PySDD (SDD), d-tree and c2d. Shaded cells contain the best times.

	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6	Q_7	Q_8	Q_9	Q_{10}	Q_{11}	Q_{12}	Q_{13}	Q_{14}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6	Q_7	Q_8	Q_9	Q_{10}	Q_{11}	Q_{12}	Q_{13}	Q_{14}
P+SDD	59	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	78	NA	150	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA
S(30)+SDD	1.3s	NA	729	NA	4.5s	817s	6s	NA	NA	NA	63	165s	30s	326	15.5	NA	8.9s	NA	NA	NA	NA	NA	NA	NA	372	NA	NA	3.3s
vP+SDD	587	7.2s	306	5.6s	13.6s	NA	6.3s	NA	NA	1.3s	2s	17.3s	12.4s	3.1s	7.3s	NA	2.5s	NA	NA	NA	NA	NA	NA	NA	2s	NA	NA	38.7s
L w/o+SDD	57	420	38	1.1s	1.3s	NA	353	35.1s	348s	187	7	10.6s	541	337	647	52s	455	2.4s	4.7s	NA	2s	51.8s	NA	1.7s	31	12.7s	6.1s	4.9s
L w/+SDD	49	383	38	175	365	NA	315	21.8s	174s	162	5	387	176	273	617	46.1s	444	1.5s	3.7s	NA	1.9s	71.4s	NA	1.6s	21	1.6s	2.8s	6s
L w/+d-tree	49	676	40	461	595	NA	4.9s	668s	108s	1.5s	6	1s	206	273	617	42s	411	1.7s	2.7s	NA	6.3s	658s	NA	2.9s	21	2s	2.4s	6s
L w/+c2d	49	41s	316	3.9s	62s	NA	27s	NA	NA	2.4s	6	13s	6.2s	273	617	NA	1s	7.4s	113s	NA	32s	NA	NA	4.2s	21	16s	16s	6s

	Query ID	2343894_40	2327997_45	2322829_40	2416754_49	2346575_46
Runtime	S(1)	1.5s	800ms	721ms	793ms	1.1s
	S(20)	1311s	148s	88s	45s	40s
	S(30)	TO	1415s	89s	42s	41s
	LTGs w/	353s	7.3s	6.1s	20s	17.6s
Probability	S(1)	0.03	0.003	0.04	0.006	0.68
	S(20)	0.12	0.02	0.05	0.007	0.97
	S(30)	TO	0.02	0.05	0.007	0.97
	LTGs w/	0.13	0.02	0.11	0.015	0.97

Table. Results over the VQAR benchmark (NeurIPS 2021).

Overview

- TGs: non-probabilistic case
- TGs: probabilistic case

Motivation

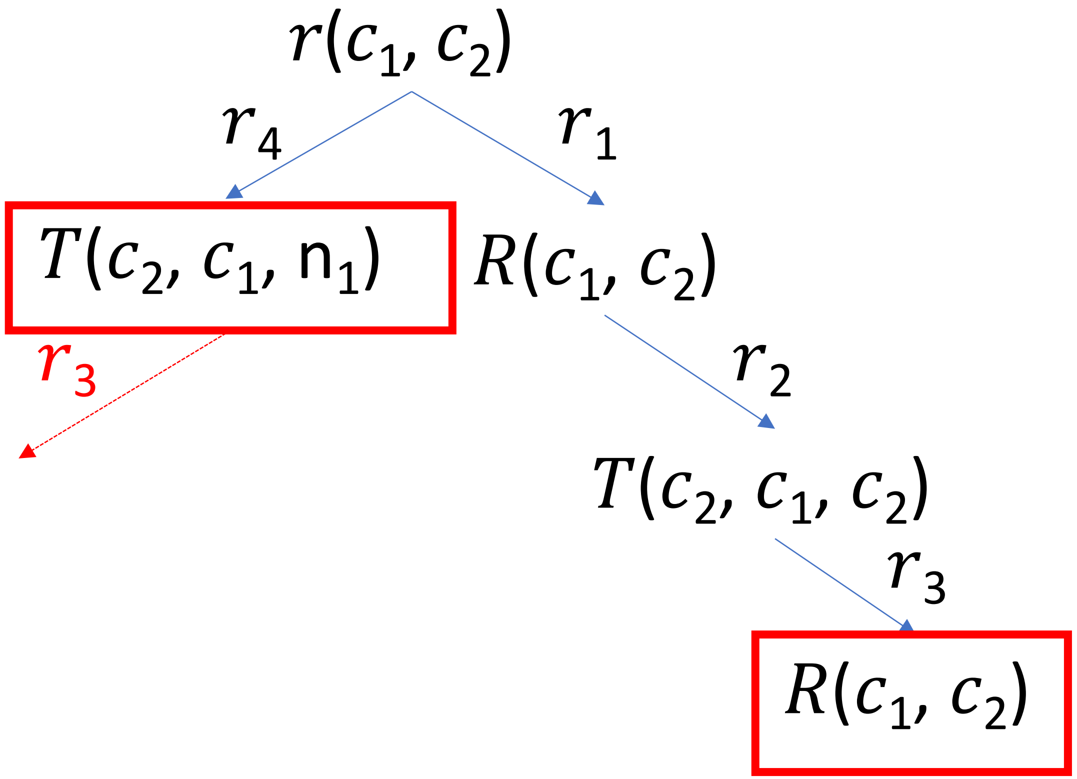

$$r(X,Y) \rightarrow R(X,Y) \quad (r_1)$$
$$R(X,Y) \rightarrow T(Y,X,Y) \quad (r_2)$$
$$T(Y,X,Y) \rightarrow R(X,Y) \quad (r_3)$$
$$r(X,Y) \rightarrow \exists Z.\, T(Y,X,Z) \quad (r_4)$$
$$B = \{r(c_1, c_2)\}$$

Standard bottom-up seminaive evaluation

- ❑ **cannot** prevent the derivation of logically redundant facts
- ❑ **cannot** prevent redundant homomorphism checks

TG Construction for Linear Rules: Intuition

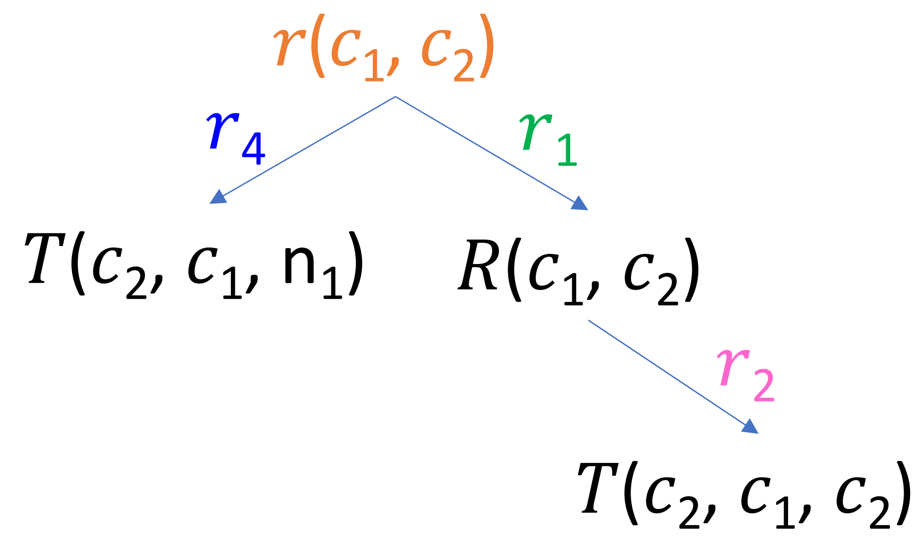

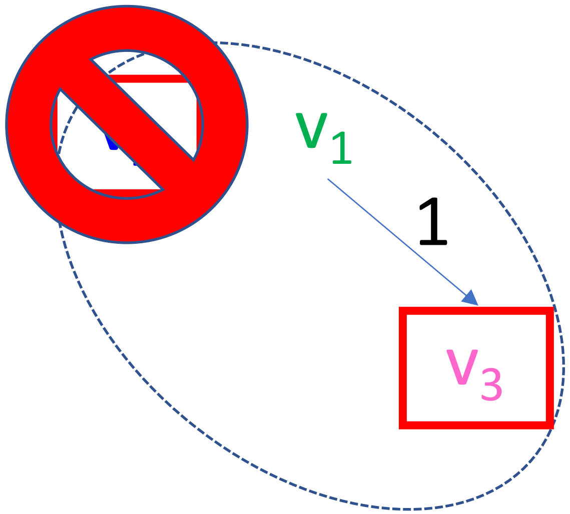

$$r(X,Y) \rightarrow R(X,Y) \quad (r_1)$$
$$R(X,Y) \rightarrow T(Y,X,Y) \quad (r_2)$$
$$T(Y,X,Y) \rightarrow R(X,Y) \quad (r_3)$$
$$r(X,Y) \rightarrow \exists Z.\, T(Y,X,Z) \quad (r_4)$$

- Reason over an instance **B*** that captures *all* possible rule execution patterns.
- Build a TG that captures the derivations over **B***.
- Eliminate nodes producing logically redundant facts:
  - preserving homomorphisms

TG-Based Reasoning: Datalog Rules

- No time to cover. Read our VLDB 2021 paper!

TG minimization

- ❑Computing rewritings over the TG.
- ❑Reduce to query containment (de cidable as the check does not con sider the rules).

Theoretical results:

- ❑Produce a minimal (all instance g uarantees) TG.
- ❑Decision problem is co-NP-compl ete.

TG-based rule execution strategy

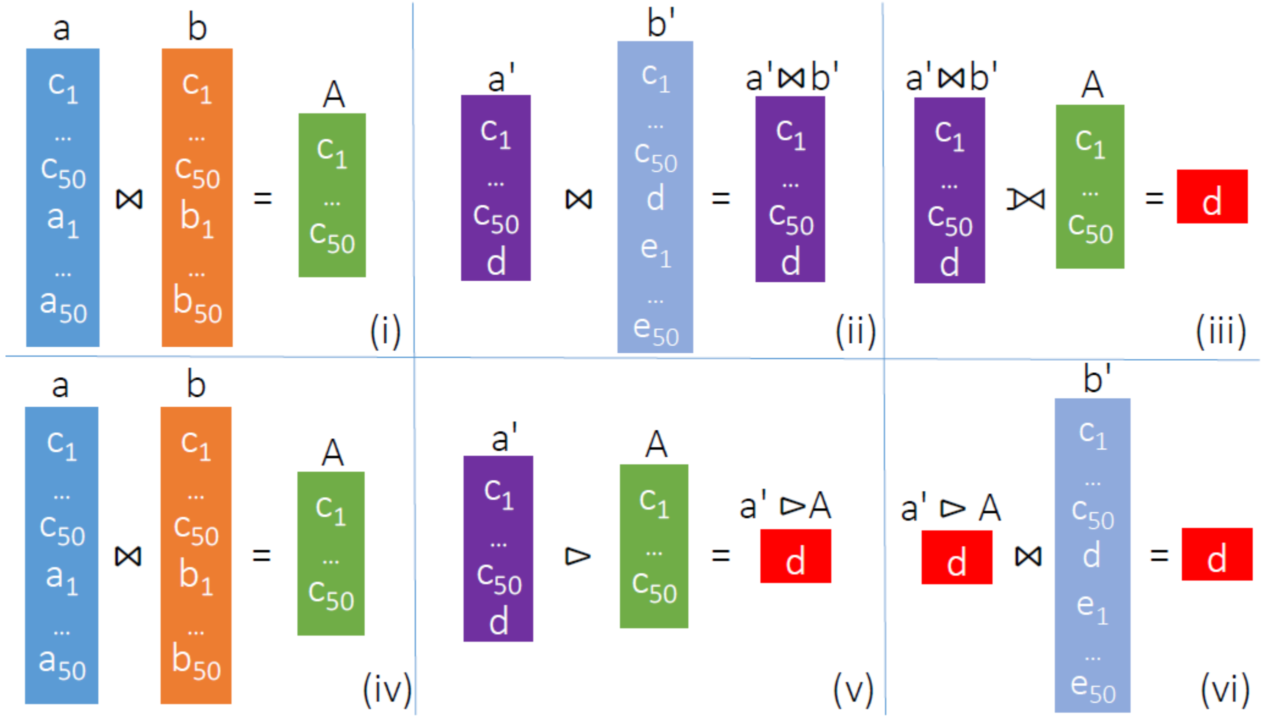

Overview

- TGs: non-probabilistic case
- TGs: probabilistic case (actually Datalog reasoning over tuple-independent PDBs)

Probabilities into the game

- Auto-mined KGs
  - Google's Knowledge Vault
  - Microsoft's Concept Graph
- Visual Question Answering

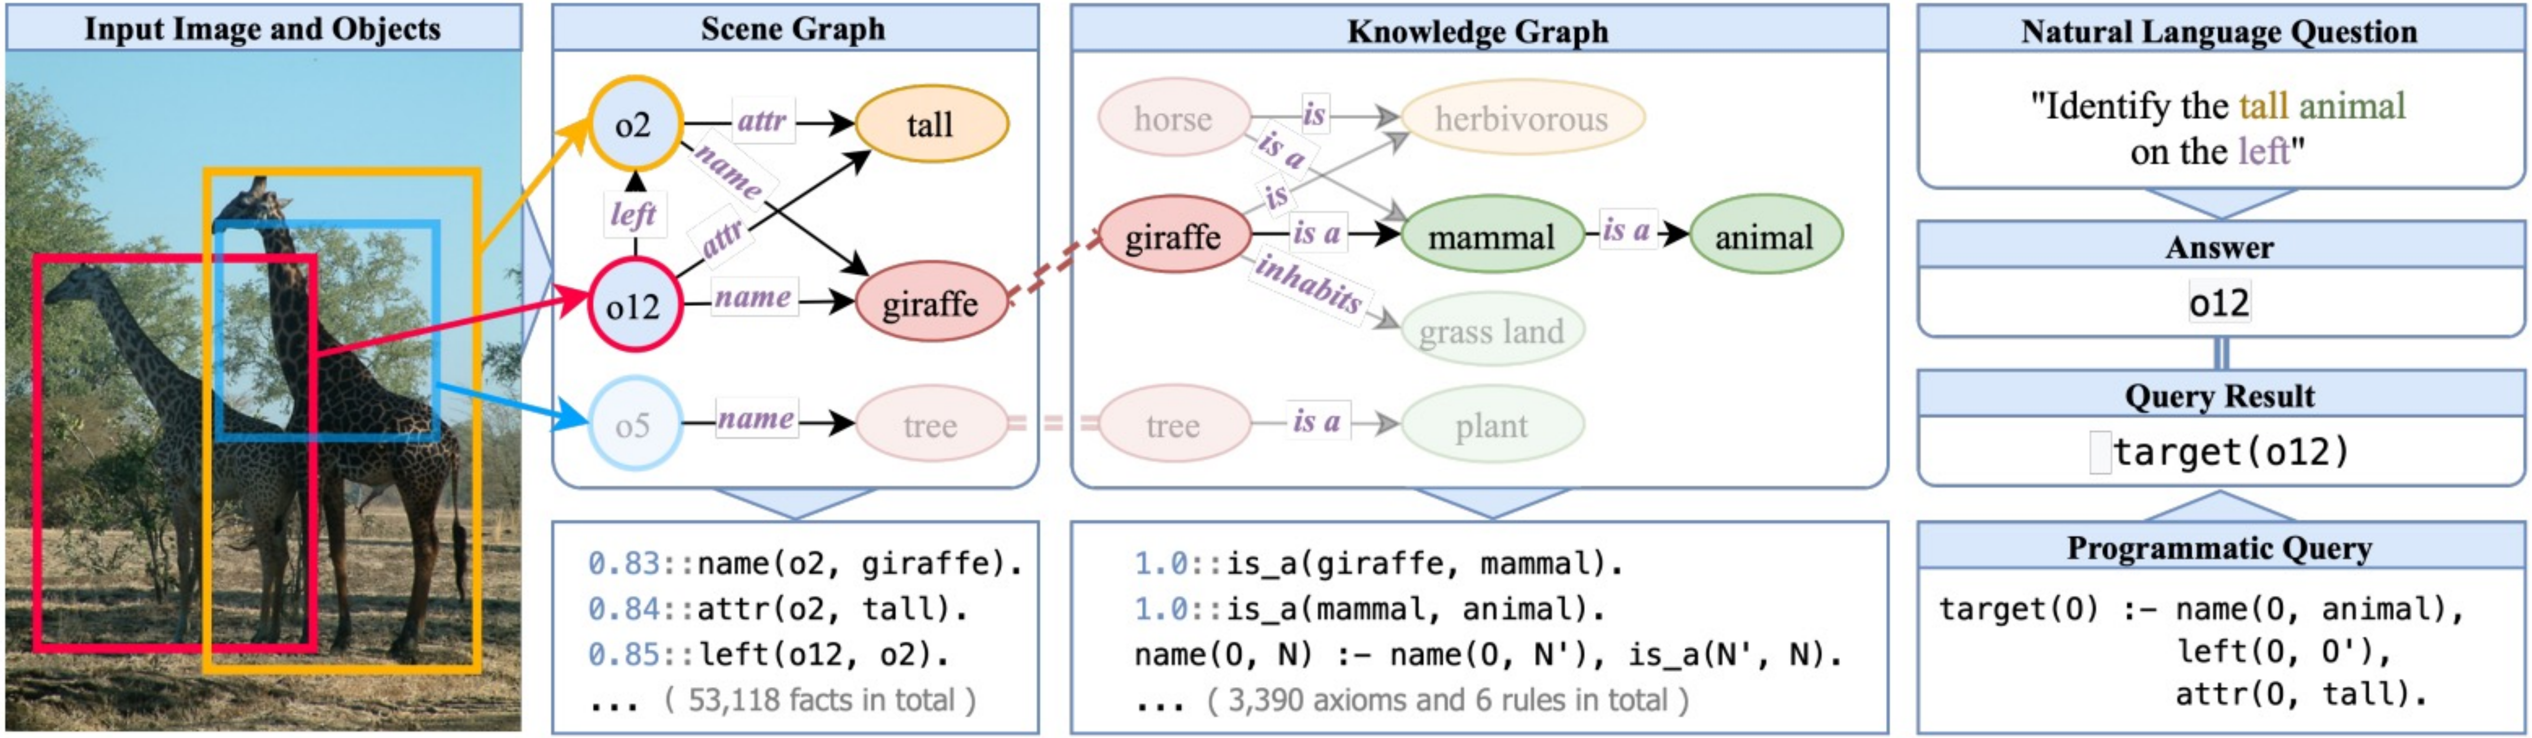

TG extensions to account for probabilities

- Why TGs should be extended?
  - We need to account for **all** possible non-redundant ways to derive each fact
  - Then, the derivations are compiled in a formula to compute the probability a derived fact is true.

TG-Based Probabilistic Reasoning: Intuition

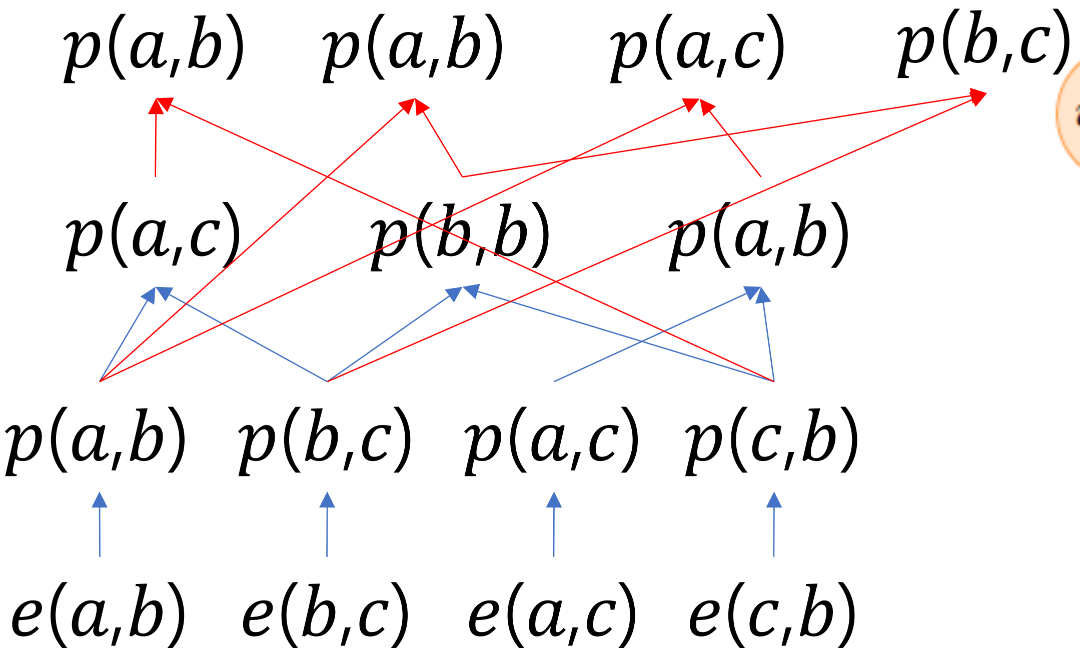

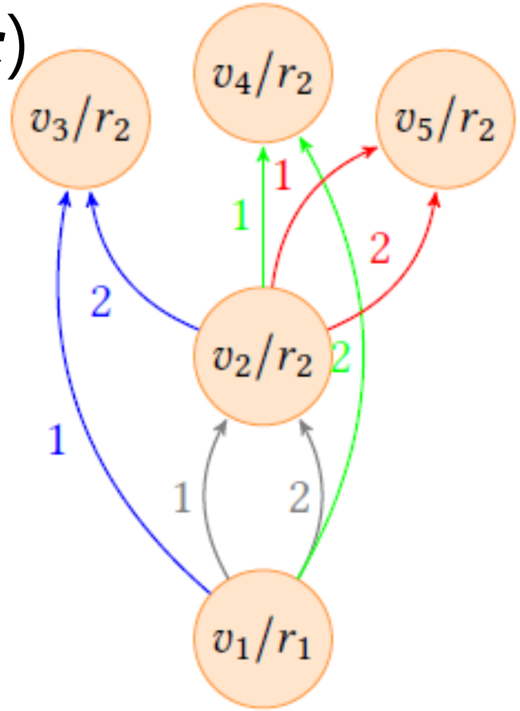

- Keep the provenance at reasoning-time within the nodes of the TG (this can be done efficiently).
- Stop when the derivation of a fact depends on itself.

$$p(X,Y) \leftarrow e(X,Y) \qquad (r_1)$$

$$p(X,Y) \leftarrow p(X,Z) \wedge p(Z,Y) \qquad (r_2)$$

$$e(a,b),\ e(b,c),\ e(a,c),\ e(c,b)$$

TG-Based Probabilistic Reasoning: Collapsing

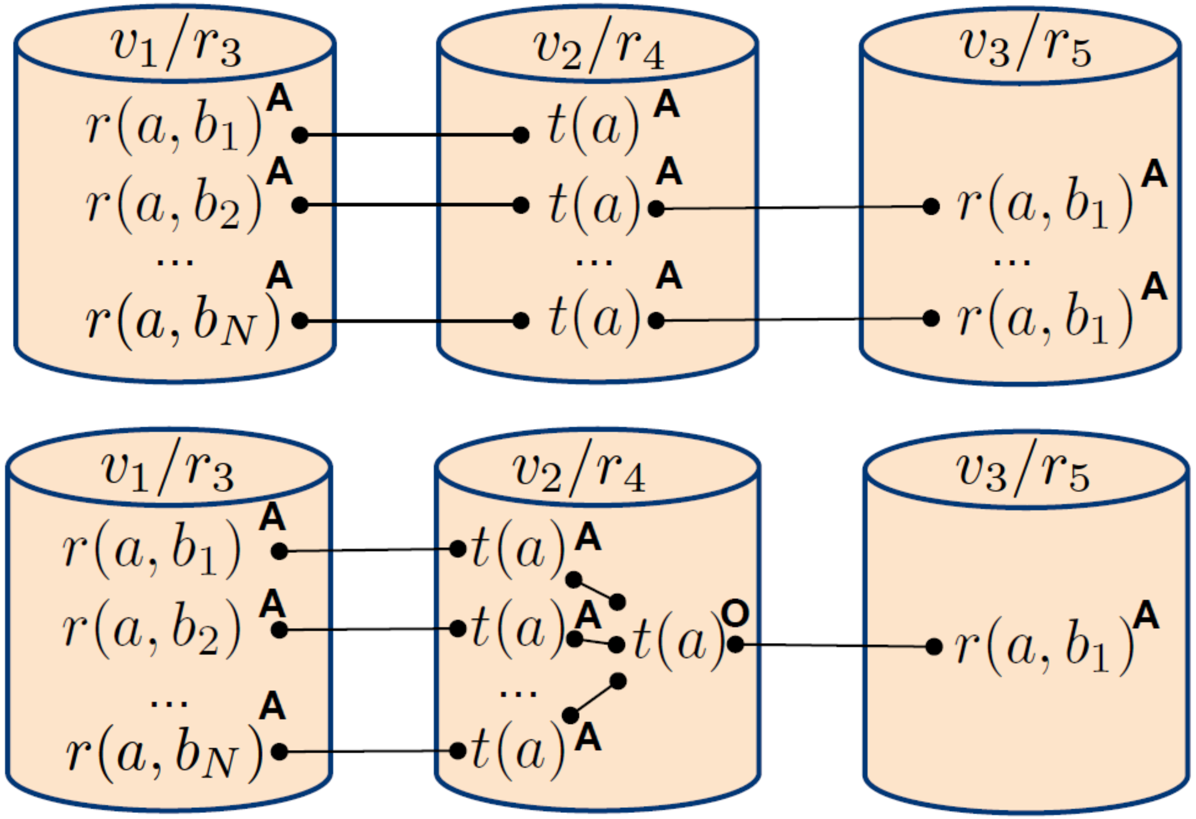

$r(X,Y) \leftarrow q(X,Y)$ (r_3)

$t(X) \leftarrow r(X,Y)$ (r_4)

$r(X,Y) \leftarrow t(X) \wedge s(X,Y)$ (r_5)

$q(a,b_i)$, for $1 \leq i \leq N$, and $s(a,b_1)$

- The previous technique is sound… however, it can explode space-wise.
- Keep only one derivation per fact within each node.
- Extends the notion of absorptive provenance circuits [D. Deutch, ICDT 2014], but we decide when to collapse or not.

Probabilistic reasoning: prior art

- How prior art [Tsamoura et al., AAAI 2020] works? Read our SIGMOD 2023 where we explain via an example!
  - They are based on **provenance semirings** [T.J. Green, PODS 2007] and have **exponential space complexity** [D. Deutch, ICDT 2014].

Connections to incremental reasoning

Samsung Research

- ???

A few last things about my research

Scene graph generation

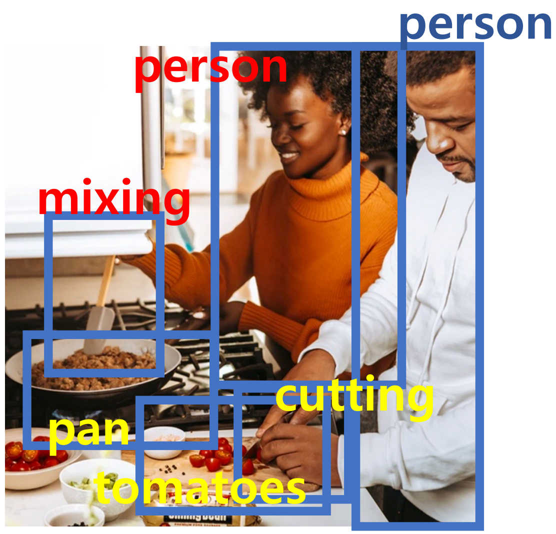

AAAI 2023

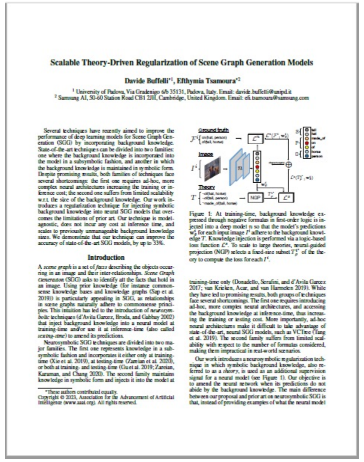

Scalable Theory-Driven Regularization of Scene Graph Generation Models

Davide Buffelli*[1], Efthymia Tsamoura*[2]

[1] University of Padova, Via Gradenigo 6/b 35131, Padova, Italy. Email: davide.buffelli@unipd.it
[2] Samsung AI, 50-60 Station Road CB1 2JH, Cambridge, United Kingdom. Email: efi.tsamoura@samsung.com

AAAI 2023

Principled and Efficient Motif Finding for Structure Learning in Lifted Graphical Models

Jonathan Feldstein [1,2], Efthymia Tsamoura[3]

[1] University of Edinburgh
[2] Bennu.ai
[3] Samsung AI
jonathan@bennu.ai, dominic.phillips@ed.ac.uk, efi.tsamoura@samsung.com

Structural motifs mining for lifted graphical models under (ε,α)-guarantees ...and some nice complexity results.

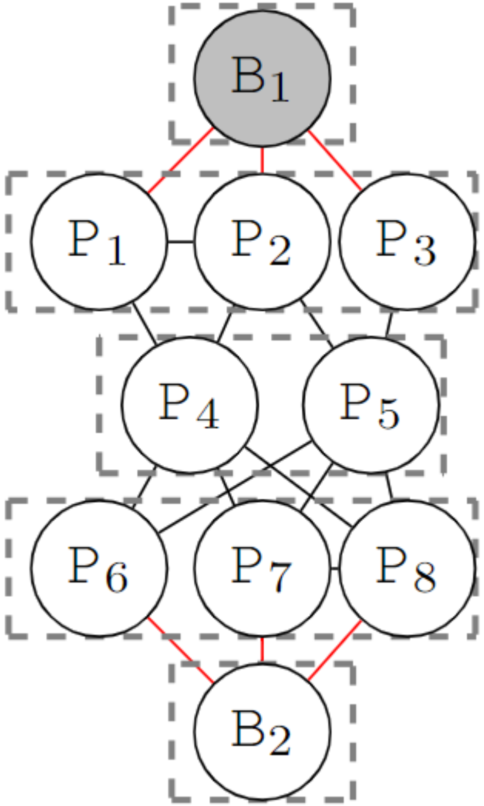

Thanks!

Samsung Research

- Please feel free to reach out if
  - ❑ you want to visit me in Cambridge, UK, or
  - ❑ you have a nice idea to work on, or
  - ❑ you want to learn more about my projects including TGs!