
Materializing Knowledge Bases via Trigger Graphs

Efi Tsamoura

Samsung AI, Cambridge, UK

About this talk

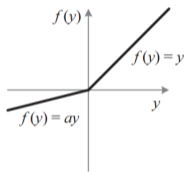
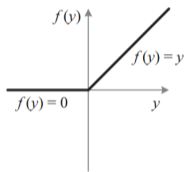
The aim is to:

- **Part I.** Motivate the use of logic and symbolic knowledge representation and reasoning techniques in developing AI applications.
- **Part II.** Present techniques to improve the efficiency of logical inference.

Part I. Motivation

Deep learning -the successes

PReLU-nets surpass humans on classification in 2015



- PReLU networks achieve **4.94%** top-5 test error on ImageNet 2012 classification.
- Human-Level top-5 test error is **5.1%**.

Kaiming He Xiangyu Zhang Shaoqing Ren Jian Sun. **Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification.** In ICCV, pages 1026–1034, 2015.

AlphaGo seals 4-1 victory over Go grandmaster in 2016

DeepMind's artificial intelligence astonishes fans to defeat human opponent and offers evidence computer software has mastered a major challenge



▶ The world's top Go player, Lee Sedol, lost the final game of the Google DeepMind challenge match. Photograph: Yonhap/Reuters

ChatGPT passes medical and law exams in 2023

ChatGPT passes exams from law and business schools



By [Samantha Murphy Kelly](#), CNN Business

🕒 4 minute read · Updated 1:35 PM EST, Thu January 26, 2023



[HOME](#) > [HEALTH](#)

The newest version of ChatGPT passed the US medical licensing exam with flying colors — and diagnosed a 1 in 100,000 condition in seconds

[Hilary Brueck](#) Apr 6, 2023, 9:03 PM BST



Deep learning strengths

- Pattern classification (in large).
- Learning via example.
- Tolerance to noise.

Deep learning -the failures

Deep (vision) models are prone to biases

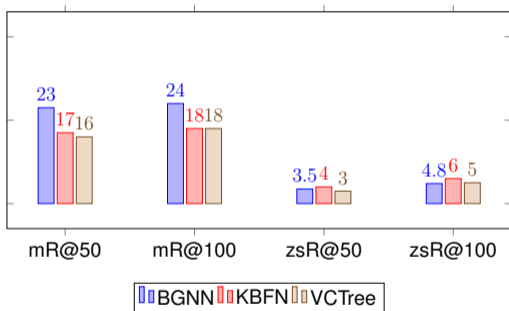
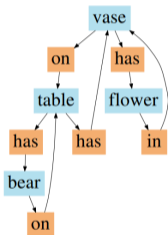
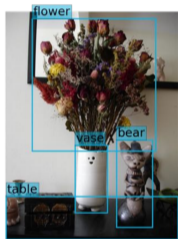


Figure: Accuracy of BGNN [11], KBFN [6] and VCTree [13] on Visual Genome [9]. Task: scene graph generation.

Large language models fail on abstract reasoning



- LLMs have very limited performance in abstract reasoning.
- Techniques that can improve performance on other NLP tasks cannot improve the abstract reasoning capabilities of large language models.

Gaël Gendron, Qiming Bao, Michael Witbrock, Gillian Dobbie. **Large Language Models Are Not Strong Abstract Reasoners.** In Arxiv, 2023.

Deep learning weaknesses

- Focus on single cognitive abilities.
- Requires large amounts of training data.
- Lacks transparency/interpretability.
- Its answers cannot be fully trusted.
- Prone to data biases.
- Difficult to incorporate background knowledge.

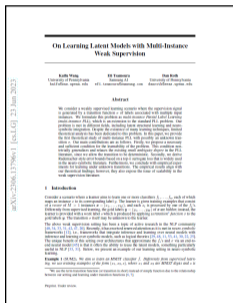


Gary Marcus. **Deep Learning: a Critical Appraisal**. In Arxiv, 2018.

Logic to the rescue

- Focus on single **complex** cognitive abilities.
- Requires large **small** amounts of training data.
- ~~Lacks~~ transparency/interpretability.
- Its answers can ~~not~~ be fully trusted.
- **Not** prone to data biases.
- ~~Difficult~~ **Straightforward** to incorporate background knowledge.

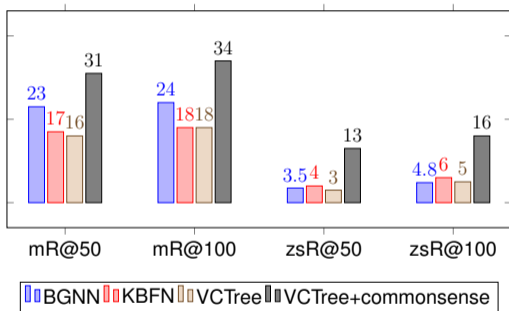
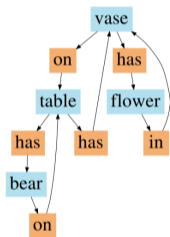
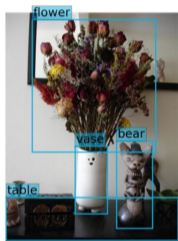
Via logic, we can indeed learn (deep) classifiers!



- (Deep) classifier learnability under **unknown** logical theories.
- (Deep) classifier learnability under **probabilistic** logical theories.

Kaifu Wang, and **Efthymia Tsamoura**, Dan Roth. **On learning latent models with multi-instance weak supervision**. In NeurIPS, 2023.

Via logic, we can overcome data biases!



Davide Buffelli, and **Efthymia Tsamoura**. **Scalable Theory-Driven Regularization of Scene Graph Generation Models**. In AAAI, 2023.

Scene Graph Generation (AAAI 2023)

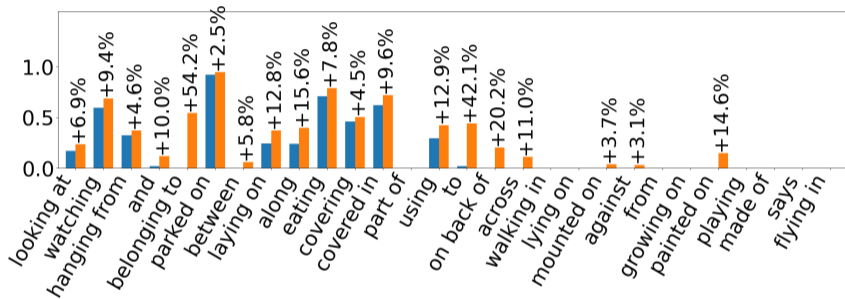
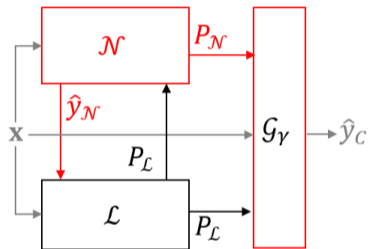


Figure: Recall of VCTree [13] on the 28 least frequent predicates: without NGP; with NGP. Benchmark: Visual Genome [9].

Knowledge Distillation into Deep Networks (ICML 2023)

Concordia



- First to support general first-order theories.
- Supports semi-/un-/supervised learning.

Operation	Equation
-----------	----------

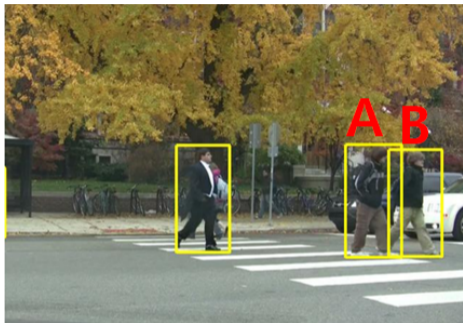
Inference	$\hat{\mathbf{y}} = \arg \max_{\mathbf{y}} P_{\mathcal{N}}(\mathbf{Y} = \mathbf{y} \mathbf{X} = \mathbf{x}, \boldsymbol{\theta})$
-----------	---

Training	$\hat{\boldsymbol{\theta}}_{t+1} = \arg \min_{\boldsymbol{\theta}} (\ell(\hat{\mathbf{y}}_{\mathcal{N}}, \mathbf{y}) + KL(P_{\mathcal{N}}, P_{\mathcal{L}}))$
----------	---

$$\hat{\boldsymbol{\lambda}}_{t+1} = \arg \max_{\boldsymbol{\lambda}} \prod_{(\mathbf{x}) \in \mathcal{D}} P_{\mathcal{L}}(\mathbf{X} = \mathbf{x}, \boldsymbol{\lambda}_t)$$

Leon Jonathan Feldstein, Jurčius Modestas and **Efthymia Tsamoura**. **Parallel neurosymbolic integration with Concordia**. In ICML (to appear), 2023.

Video Activity Detection (ICML 2023)



$SEQ(B_1, B_2) \wedge CLOSE(B_1, B_2) \rightarrow SAME(B_1, B_2)$

$DOING(B_1, A) \wedge SAME(B_1, B_2) \rightarrow DOING(B_2, A)$

Accuracy over 5 runs

Model	Avg (%)	Max (%)	Min (%)
ACD+ \mathcal{L} [12]	86.00	-	-
MobileNet	90.07	91.36	89.61
IARG(MobileNet) [10]	90.18	92.39	87.55
Concordia (MobileNet, \mathcal{L})	90.73	93.19	89.54
Inception	89.72	91.83	86.84
IARG(Inception) [10]	88.88	91.67	85.33
Concordia (Inception, \mathcal{L})	92.75	93.34	92.31

Leon Jonathan Feldstein, Jurčius Modestas and **Efthymia Tsamoura**. **Parallel neurosymbolic integration with Concordia**. In ICML (*to appear*), 2023.

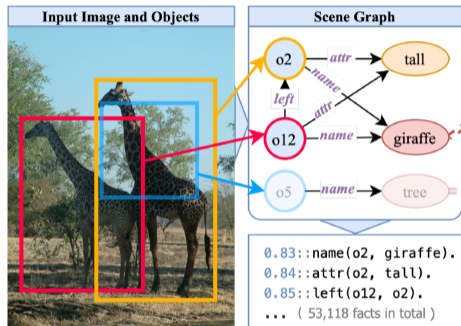
Entity Linking (ICML 2023)

Table: Results on entity linking.

Model	F_1	Acc (%)
BERT (sp)	0.88	88.5
Concordia(BERT) (sm)	0.91	91.4

Leon Jonathan Feldstein, Jurčius Modestas and **Efthymia Tsamoura**. **Parallel neurosymbolic integration with Concordia**. In ICML, 2023 (*to appear*).

Visual QA (SIGMOD 2023)



$Q(O) \leftarrow \text{NAME}(\textit{herbivore}, O)$
 $\text{NAME}(N, O) \wedge \text{NAME}(N', O) \rightarrow \text{ISA}(N', N)$
 $\rightarrow \text{ISA}(\textit{giraffe}, \textit{herbivore})$
 $\rightarrow \text{ISA}(\textit{deer}, \textit{herbivore})$

Table: Recall@5 on VQAR [7].

Testset	LXMERT [14]	RVC [4]	TG-Guided VQA
C5	64.05%	74.62%	87.01%
C6	56.51%	72.04%	85.45%

Efthymia Tsamoura, Jaehun Lee, and Jacopo Urbani. **Probabilistic Reasoning as Scale: Trigger Graphs to the Rescue.** In SIGMOD, 2023 (to appear).

Part II: Reasoning at Scale via Trigger Graphs

Efthymia Tsamoura, David Carral, Enrico Malizia, and Jacopo Urbani. **Materializing Knowledge Bases via Trigger Graphs**. In VLDB, pages 943-951, 2021.

Trigger Graphs: Why

- **Key to support goal-driven QA over transitive rules.**
- **Standard bottom-up evaluation:**
 - may derive logically redundant facts;
 - may try to execute rules that derive no facts.
- **The above negatively impact the runtime and the memory.**

How: Trigger Graphs

Rules

$$r(X, Y) \rightarrow R(X, Y) \quad (r_1)$$

$$R(X, Y) \rightarrow T(Y, X, Y) \quad (r_2)$$

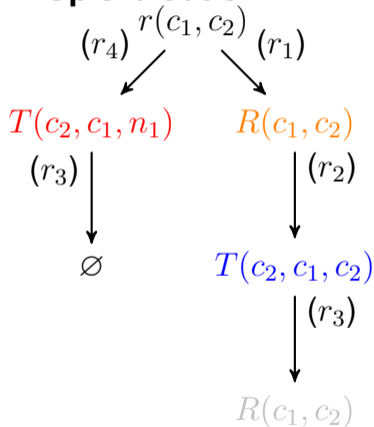
$$T(Y, X, Y) \rightarrow R(X, Y) \quad (r_3)$$

$$r(X, Y) \rightarrow \exists Z.T(Y, X, Z) \quad (r_4)$$

Facts

$$\rightarrow r(c_1, c_2)$$

Bottom-Up evaluation



How: Trigger Graphs

Rules

$$r(X, Y) \rightarrow R(X, Y) \quad (r_1)$$

$$R(X, Y) \rightarrow T(Y, X, Y) \quad (r_2)$$

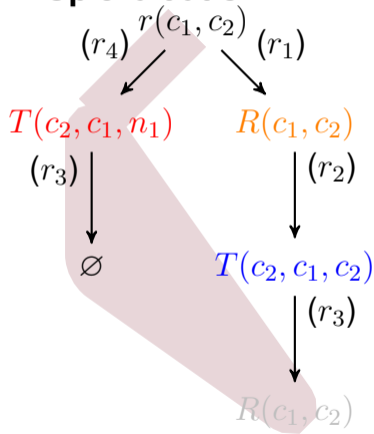
$$T(Y, X, Y) \rightarrow R(X, Y) \quad (r_3)$$

$$r(X, Y) \rightarrow \exists Z.T(Y, X, Z) \quad (r_4)$$

Facts

$$\rightarrow r(c_1, c_2)$$

Bottom-Up evaluation



How: Trigger Graphs

Rules

$$r(X, Y) \rightarrow R(X, Y) \quad (r_1)$$

$$R(X, Y) \rightarrow T(Y, X, Y) \quad (r_2)$$

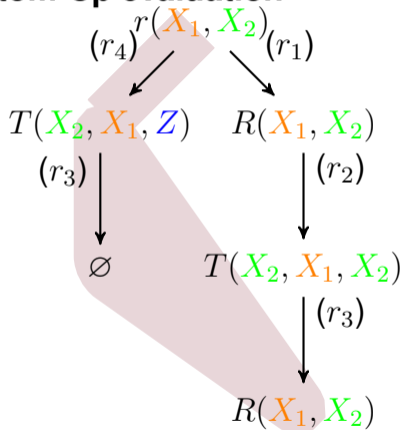
$$T(Y, X, Y) \rightarrow R(X, Y) \quad (r_3)$$

$$r(X, Y) \rightarrow \exists Z. T(Y, X, Z) \quad (r_4)$$

Facts

$$\rightarrow r(c_1, c_2)$$

Bottom-Up evaluation



How: Trigger Graphs

Rules

$$r(X, Y) \rightarrow R(X, Y) \quad (r_1)$$

$$R(X, Y) \rightarrow T(Y, X, Y) \quad (r_2)$$

$$T(Y, X, Y) \rightarrow R(X, Y) \quad (r_3)$$

$$r(X, Y) \rightarrow \exists Z. T(Y, X, Z) \quad (r_4)$$

Facts

$$\rightarrow r(c_1, c_2)$$

Trigger graph



Trigger graph-based reasoning

TGs delineate the rule executions

- Execute r_1 over the input instance.
- Execute r_2 over the derivations of r_1 .
- **No other operation is taking place.**

Important to node

- Facts are stored inside the nodes, i.e., not stored in a single set like in all bottom-up engines.
- This data separation makes joins run faster.

Trigger graph



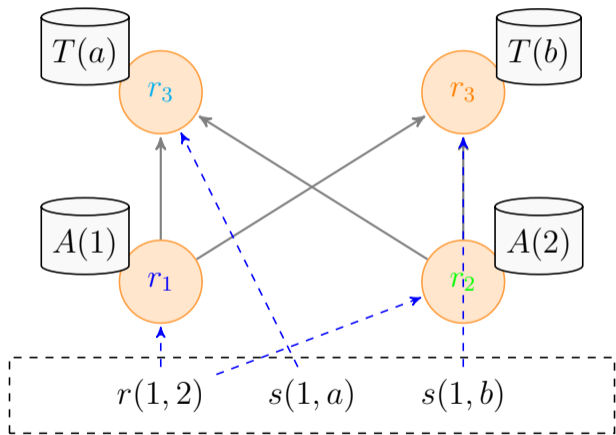
Trigger graph-based reasoning

Rules

$$r(X, Y) \rightarrow A(X) \quad (r_1)$$

$$r(X, Y) \rightarrow A(Y) \quad (r_2)$$

$$A(X) \wedge s(X, Z) \rightarrow T(Z) \quad (r_3)$$



Trigger Graphs for Linear Rules

- **Phase I: Static TG Computation.**
 - Compute a *representative* instance B^* , i.e., one that captures *all* possible rule execution paths.
 - Compute a *plan* G that mimics the rule execution when reasoning over B^* .
- **Phase II: Redundancy Elimination.**
 - Eliminate nodes that lead to redundant facts (via detecting preserving homomorphisms).
- **Phase III: Reasoning.**
 - The computed TG can be used to reason over *all* input instances.

Trigger Graphs for Linear Rules: Complexity

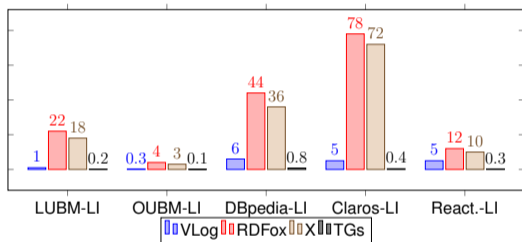
Let P be a linear program that admits a finite universal model.

Theorem (Complexity)

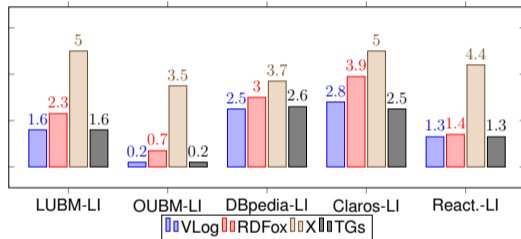
Computing a TG for P is double exponential in P . If the arity of the predicates in P is bounded, the computation time is (single) exponential.

Reasoning over Linear Rules

Total materialization times in s



Pick memory in GB



Trigger Graphs for Datalog Rules

TGs for Linear Rules

- Static TG computation.
- Use the pre-computed TG to reason over *all* instances.
- Redundancy elimination via detecting preserving homomorphisms.

TGs for Datalog Rules

- Interleave TG creation with reasoning.
- The computed TG can be used to reason over the given instance only.
- Redundancy elimination via query containment [2].

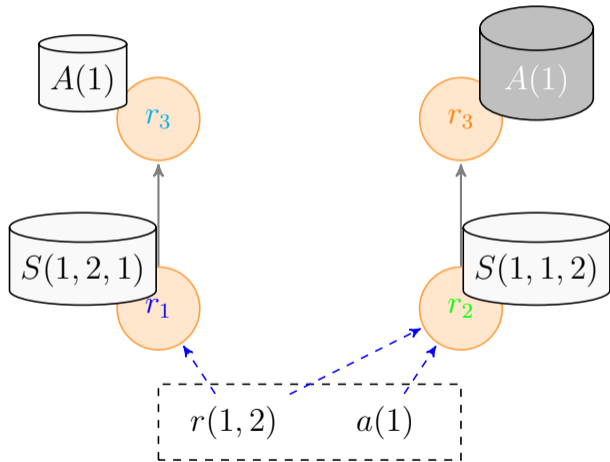
Trigger Graphs for Datalog Rules: Example

Rules

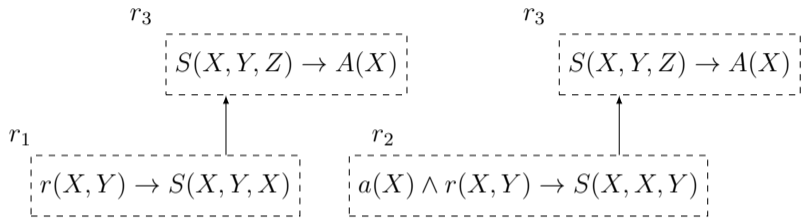
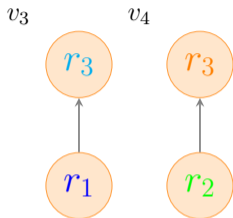
$$r(X, Y) \rightarrow S(X, Y, X) \quad (1)$$

$$a(X) \wedge r(X, Y) \rightarrow S(X, X, Y) \quad (2)$$

$$S(X, Y, Z) \rightarrow A(X) \quad (3)$$



Trigger Graphs for Datalog Rules: Example



$$Q(X) = \exists Y.r(X, Y)$$

Query for v_3

$$Q'(X) = \exists Y.a(X) \wedge r(X, Y)$$

Query for v_4

Trigger Graphs for Datalog Rules: Results

Let P be a Datalog program.

Theorem (Soundness)

For a TG G for P , $\text{minDatalog}(G)$ is a TG for P .

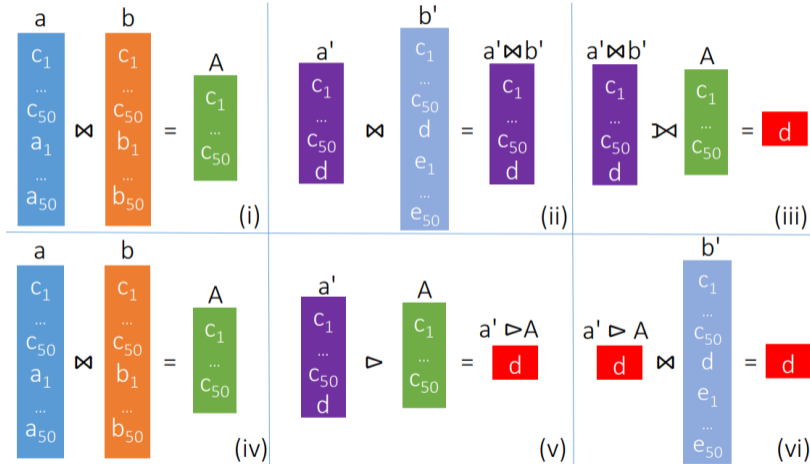
Theorem (Minimality)

Any TG for P has at least as many nodes as $\text{minDatalog}(G)$.

Theorem (Complexity)

Deciding whether G is a TG of minimum size for P is co-NP-complete.

More: TG-Aware Rule Execution Strategy



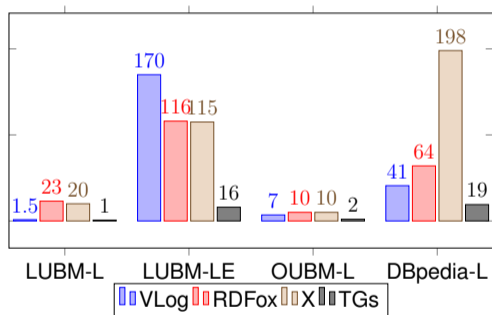
Datalog Reasoning with Trigger Graphs

	1B	2B	4B	8B	17B
Runtime (s)	203	226	520	993	2272
Memory (GB)	23	34	49	98	174
#IDPs	1B	2B	5B	10B	20B

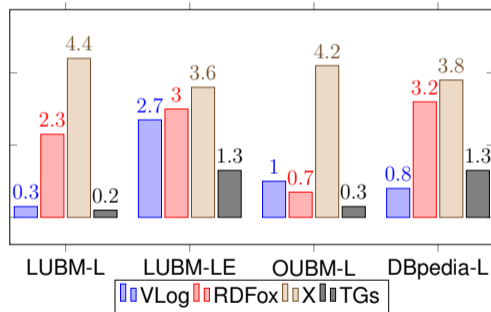
Table: Reasoning over LUBM for 1B–17B of database triples.

Datalog Reasoning with Trigger Graphs

Materialization times in s

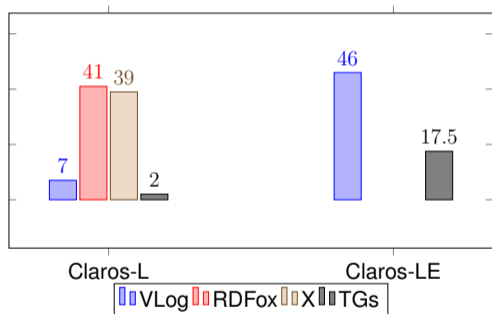


Pick memory in GB

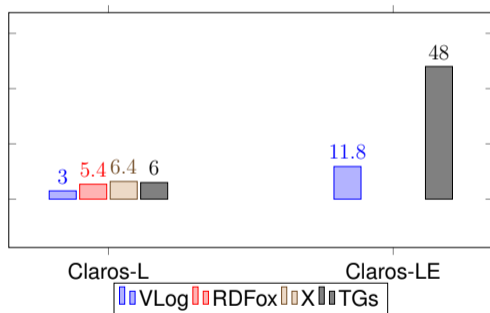


Datalog Reasoning with Trigger Graphs

Materialization times in minutes



Pick memory in GB



Reasoning at Scale: How -Lineage Trigger Graphs

Efthymia Tsamoura, Jaehun Lee, and Jacopo Urbani. **Probabilistic Reasoning as Scale: Trigger Graphs to the Rescue**. In SIGMOD, 2023 (*to appear*).

Aim

- Develop highly-scalable reasoning techniques that support uncertainty.
- Adopt well-established semantics.

Key Challenge: Complexity

Rules

$$e(X, Y) \rightarrow p(X, Y)$$

$$p(X, Z) \wedge p(Z, Y) \rightarrow p(X, Y)$$

Facts

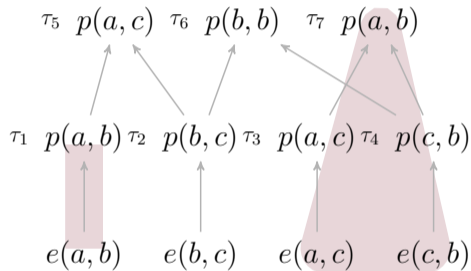
$$\rightarrow e(a, b)$$

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

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

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

Derivations



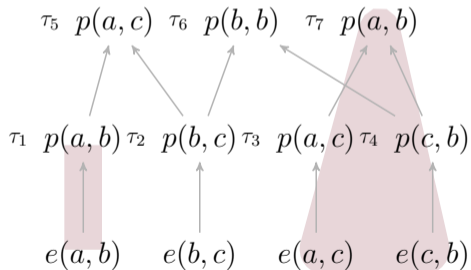
Prior Art: Key Limitations

- Relies on provenance semirings [5], i.e., associates a Boolean formula to each derivation.
- Super-polynomial size blowup in data complexity: *any monotone formula to test connectivity in a graph with n nodes has size $n^{\Omega(\log n)}$ (lower bound holds even for undirected graphs) [8].*
- Requires Boolean checks at each reasoning step for termination.
 - Runtime bottleneck.

Efthymia Tsamoura, Victor Gutierrez-Basulto, and Angelika Kimmig. **Beyond the Grounding Bottleneck: Datalog Techniques for Inference in Probabilistic Logic Programs**. In AACL, pages 10284-10291, 2020.

Probabilistic Reasoning via Provenance Semirings

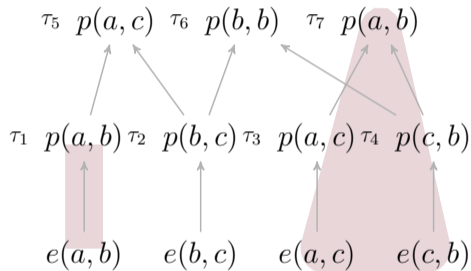
R	Derivation@R	Comparison	Formula@R
1	$e(a, b)$	\emptyset	$e(a, b)$
2	$e(a, c) \wedge e(c, b)$	$e(a, c) \wedge e(c, b) \stackrel{?}{\equiv} e(a, b)$	$e(a, c) \wedge e(c, b) \vee e(a, b)$



Lineage Trigger Graphs

- Efficient maintenance of derivation history.
- Natural for TGs.
- Storing pointer offsets.
- Reduces termination checks for detecting cyclic derivations!
 - No Boolean checks are required!

Derivations



Lineage Trigger Graphs: (Adaptive) Provenance Circuits

- Extended the notion of provenance circuits [3] to allow a more space-efficient reasoning:
 - Polynomial size representation.

Probabilistic Datalog Reasoning with Trigger Graphs

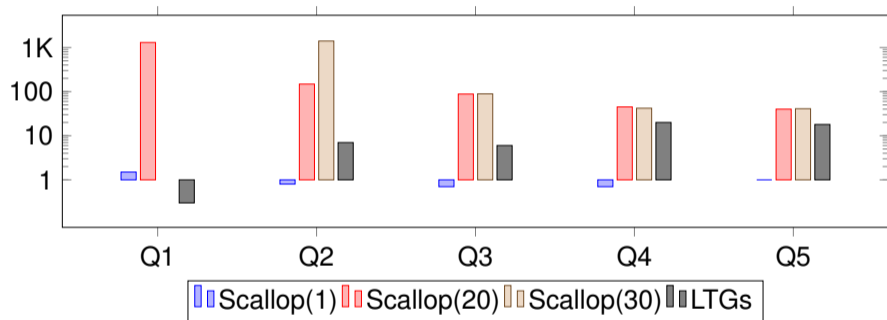


Figure: Time in seconds for goal-driven QA over sample queries from VQAR [7].

Cool Research not Covered

Goal-driven QA over existential rules with equality (AAAI 2018)

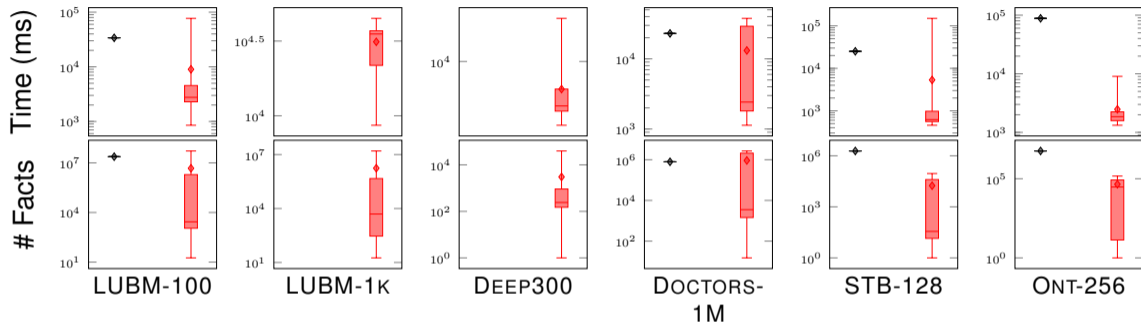


Figure: Time in msec to answer the ChaseBench queries [1].

Michael Benedikt, Boris Motik, and Efhymia Tsamoura. Goal-Driven Query Answering over Existential Rules with Equality. In AAAI, pages 1761–1770, 2018.



Thanks!

Contact info: `efi.tsamura@samsung.com`.



References I

-  Michael Benedikt, George Konstantinidis, Giansalvatore Mecca, Boris Motik, Paolo Papotti, Donatello Santoro, and Efthymia Tsamoura.
Benchmarking the Chase.
In *PODS*, pages 37–52, 2017.
-  Ashok K. Chandra and Philip M. Merlin.
Optimal implementation of conjunctive queries in relational data bases.
In *STOC*, pages 77–90, 1977.
-  Daniel Deutch, Tova Milo, Sudeepa Roy, and Val Tannen.
Circuits for datalog provenance.
In *ICDT*, pages 201–212, 2014.

References II

-  Difei Gao, Ruiping Wang, Shiguang Shan, and Xilin Chen.
From two graphs to N questions: A VQA dataset for compositional reasoning on vision and commonsense.
CoRR, abs/1908.02962, 2019.
-  Todd J. Green, Grigoris Karvounarakis, and Val Tannen.
Provenance semirings.
In *PODS*, page 31–40, 2007.
-  Jiuxiang Gu, Handong Zhao, Zhe Lin, Sheng Li, Jianfei Cai, and Mingyang Ling.
Scene graph generation with external knowledge and image reconstruction.
In *CVPR*, pages 1969–1978, 2019.

References III

-  Jiani Huang, Ziyang Li, Binghong Chen, Karan Samel, Mayur Naik, Le Song, and Xujie Si.
Scallop: From probabilistic deductive databases to scalable differentiable reasoning.
In *NeurIPS*, pages 25134–25145, 2021.
-  Mauricio Karchmer and Avi Wigderson.
Monotone circuits for connectivity require super-logarithmic depth.
In *STOC*, page 539–550, 1988.

References IV

 Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei.

Visual genome: Connecting language and vision using crowdsourced dense image annotations.

Int. J. Comput. Vis., 123(1):32–73, 2017.

 Zijian Kuang and Xinran Tie.

Video understanding based on human action and group activity recognition.

CoRR, abs/2010.12968, 2020.

References V

-  Rongjie Li, Songyang Zhang, Bo Wan, and Xuming He.
Bipartite graph network with adaptive message passing for unbiased scene graph generation.
In *CVPR*, pages 11109–11119, 2021.
-  B. London, S. Khamis, S. H. Bach, B. Huang, L. Getoor, and L. Davis.
Collective activity detection using hinge-loss markov random fields.
In *CVPR Workshops*, pages 566–571, 2013.
-  Knot Pipatsrisawat and Adnan Darwiche.
New compilation languages based on structured decomposability.
In *AAAI*, page 517–522, 2008.

References VI



Hao Tan and Mohit Bansal.

LXMERT: Learning cross-modality encoder representations from transformers.

In *EMNLP*, pages 5100–5111, 2019.