Knowledge Graph Embeddings for Biodata

Luca Costabello
Research Scientist
Accenture Labs
@lukostaz

Masaryk University
November 27, 2023

Accenture Labs BioInnovation

- Enable enhanced diagnostics and treatments, including novel therapeutics.
- Advance understanding of disease mechanisms and the effects of environmental factors.
- Deliver leaner and more effective multi-omics-driven drug discovery.

Current R&D agenda:

- AI for Pre-clinical Drug Discovery
- Biodata and ML for Genomic Medicine
- AI for Synthetic Biology

Core technologies: Biodata, Machine Learning, Knowledge Graphs, & Graph ML, Explainable AI, LLMs

Knowledge Graphs

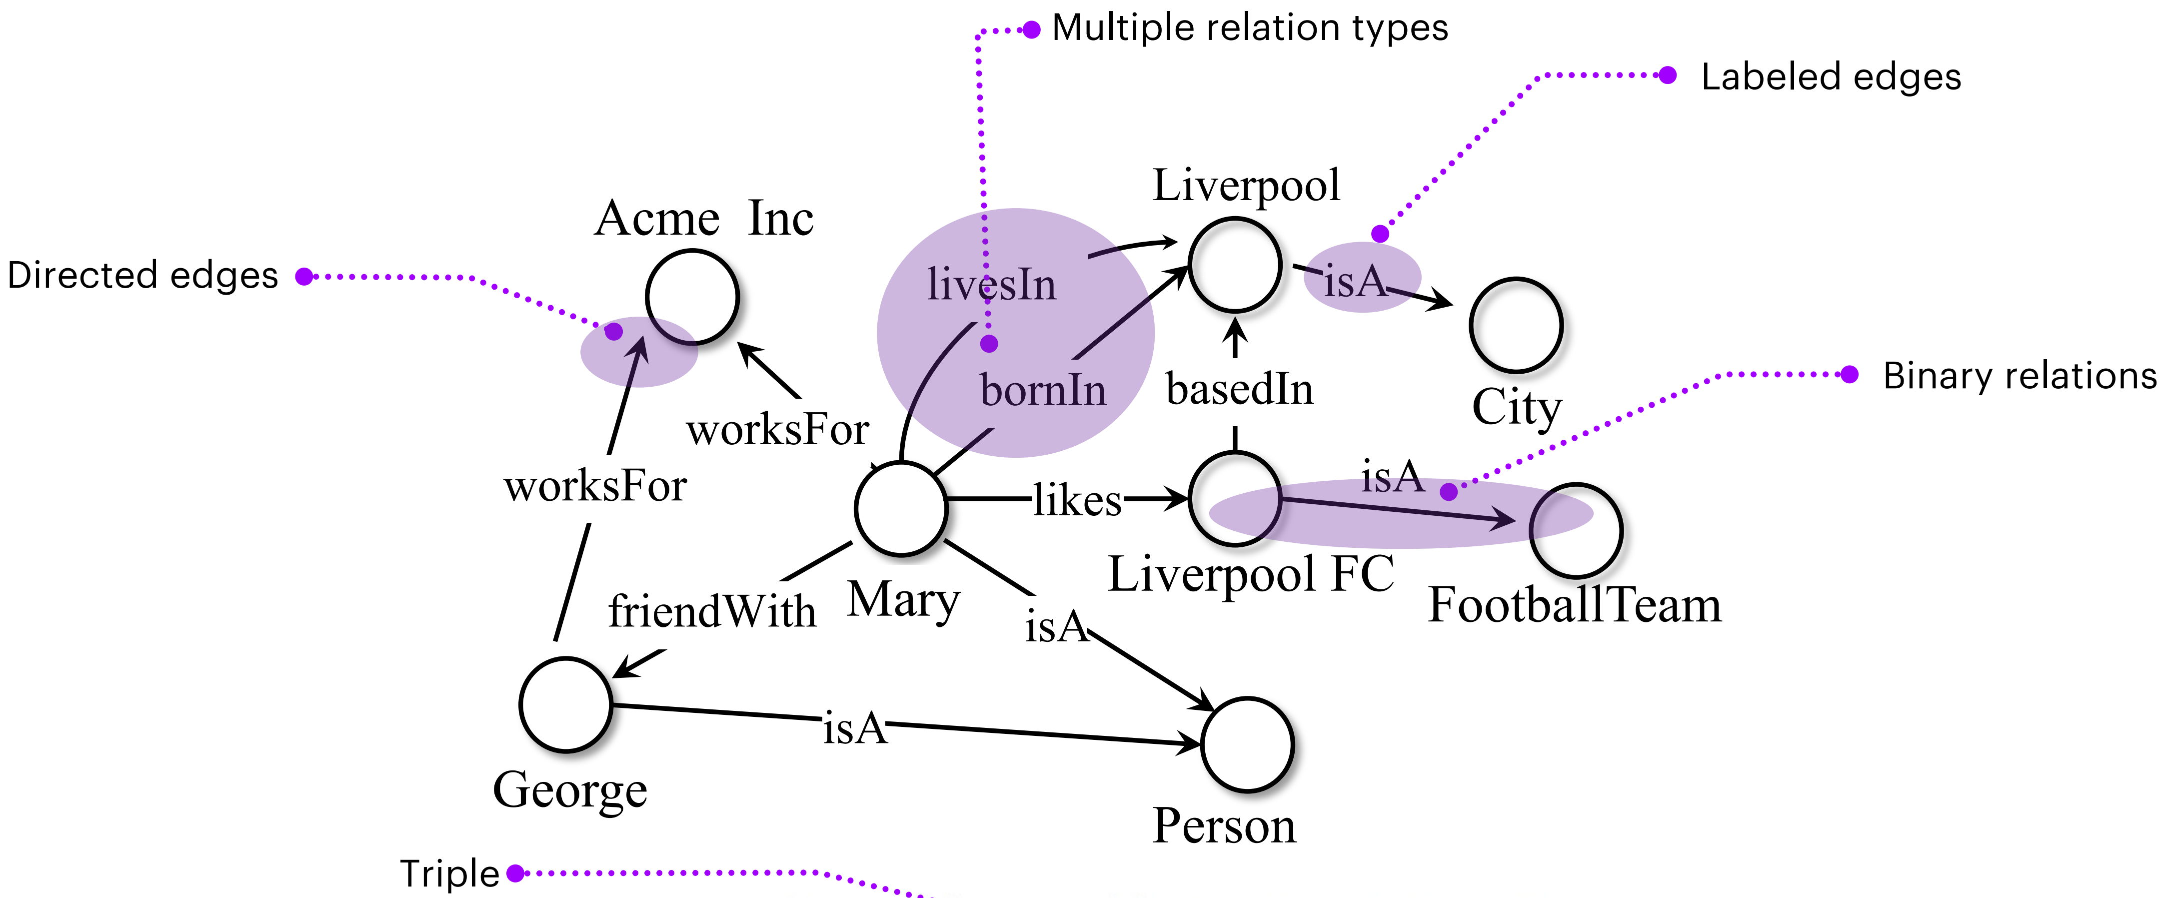

$$\mathcal{G} = \{(s, p, o)\} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$$

$\mathcal{E}$: set of entities of $\mathcal{G}$ $\qquad$ $\mathcal{R}$: set of relations of $\mathcal{G}$

In-depth overview of Knowledge Graphs in [Hogan et al. 2020]

Knowledge Graph Embeddings for Biodata

Luca Costabello
Research Scientist
Accenture Labs
@lukostaz

Trinity College Dublin
October 16, 2023

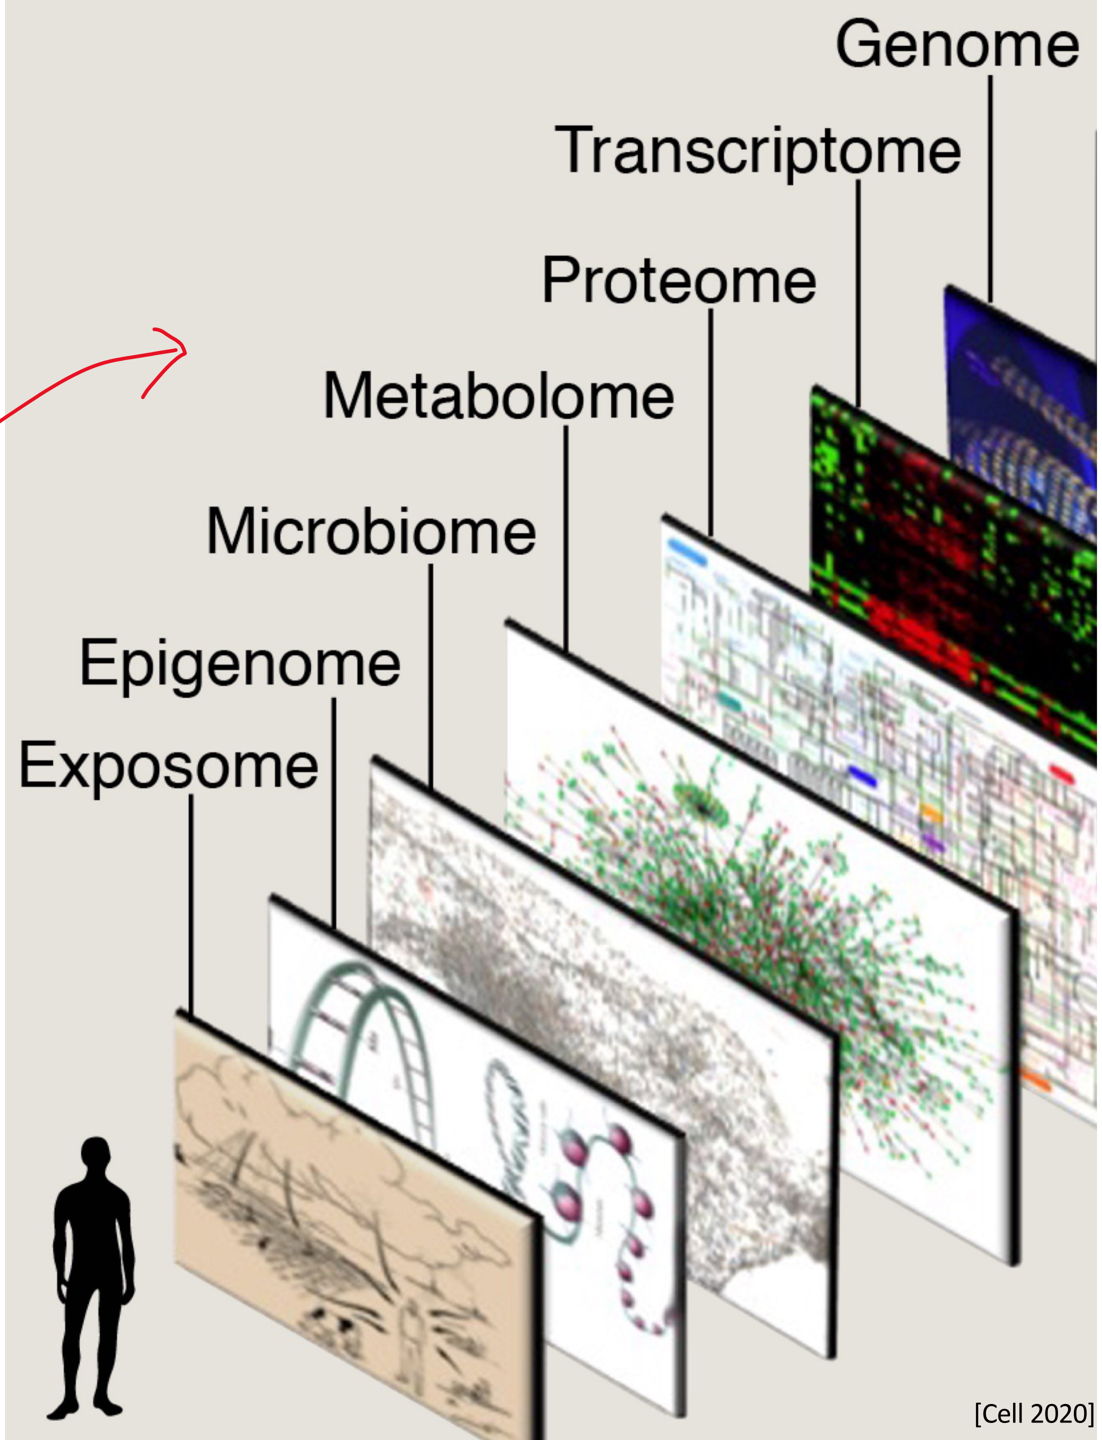

[Cell 2020]

Knowledge Graph Embeddings for Biodata

Luca Costabello
Research Scientist
Accenture Labs
@lukostaz

Trinity College Dublin
October 16, 2023

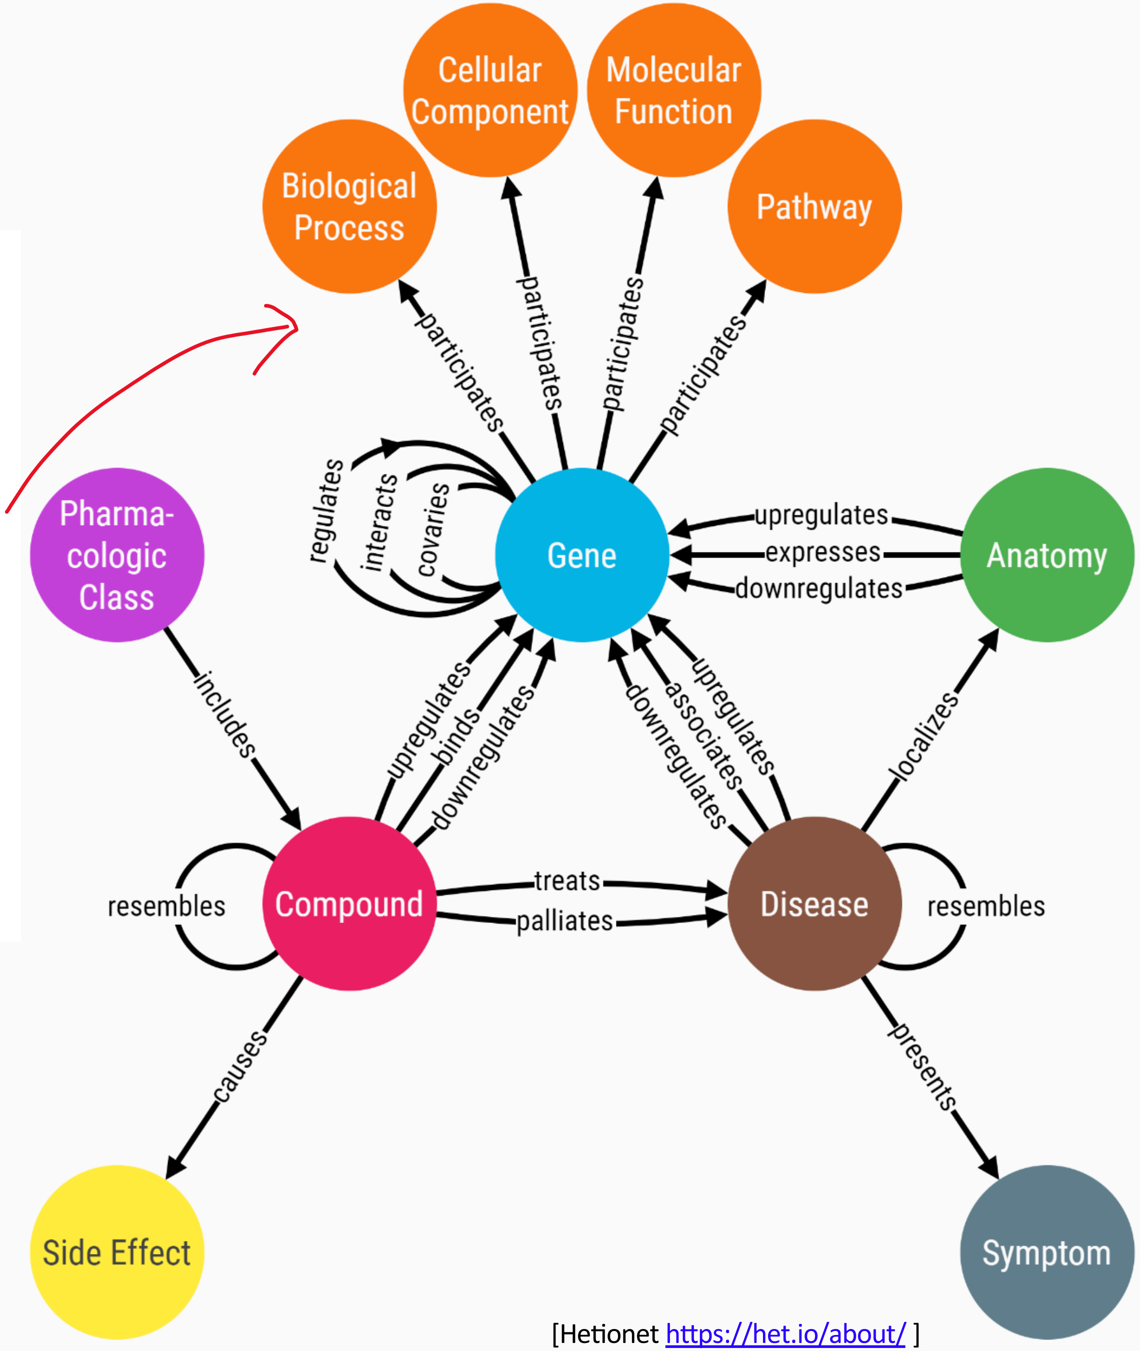

[Hetionet https://het.io/about/]

Machine Learning on Knowledge Graphs: Tasks

LINK PREDICTION / TRIPLE CLASSIFICATION

- Knowledge graph completion
- Content recommendation
- Knowledge discovery

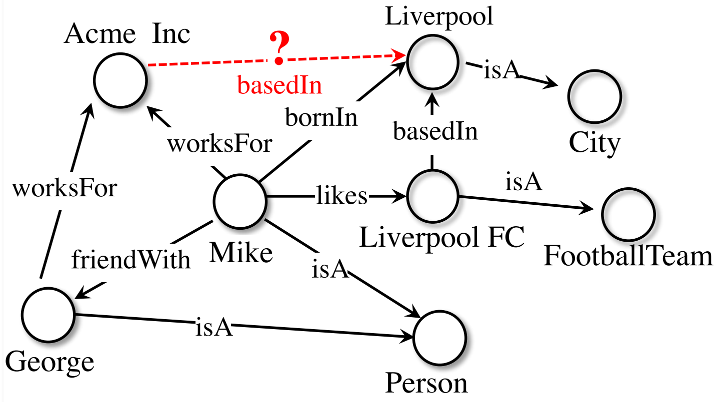

COLLECTIVE NODE CLASSIFICATION / LINK-BASED CLUSTERING

- Customer segmentation

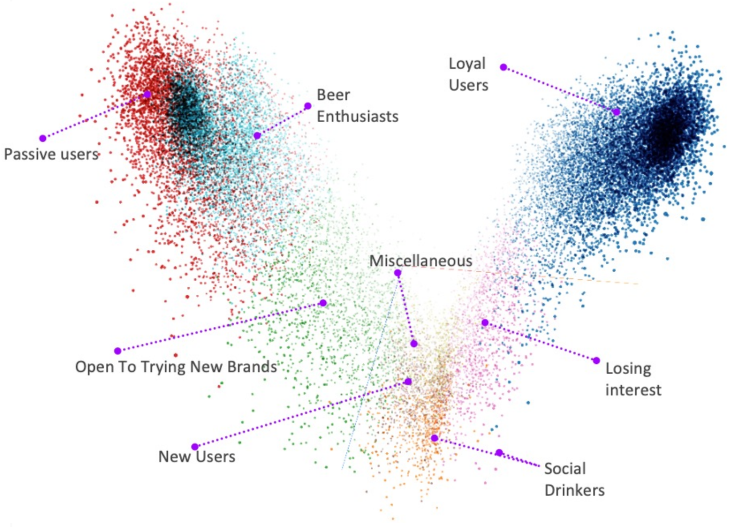

[Pai et al. 2022]

ENTITY MATCHING

- Duplicate detection
- Inventory items deduplication

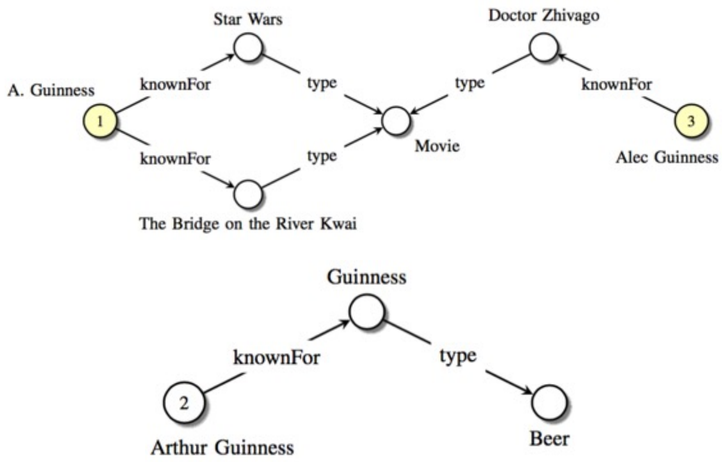

[Nickel et al. 2016a]

LINK PREDICTION / TRIPLE CLASSIFICATION

- Knowledge graph completion
- Content recommendation
- Question answering

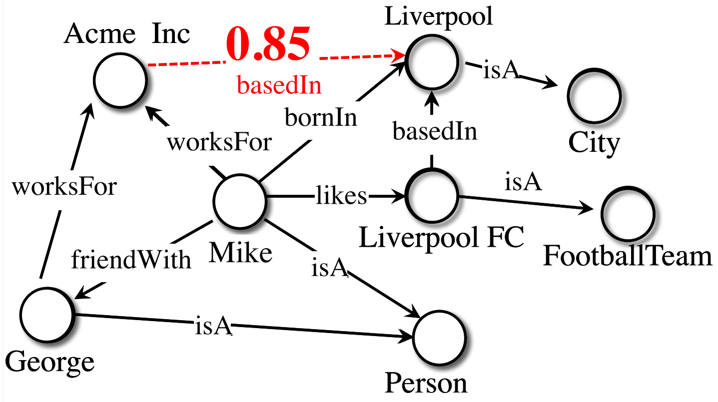

Assigning a score proportional to the likelihood that an unseen triple is true.

Link Prediction

- Learning to rank problem
- Information retrieval metrics
- No ground truth negatives in test set required

Triple Classification

- Binary classification task
- Binary classification metrics
- Test set requires positives and ground truth negatives

Graph Representation Learning

Learning representations of nodes and edges

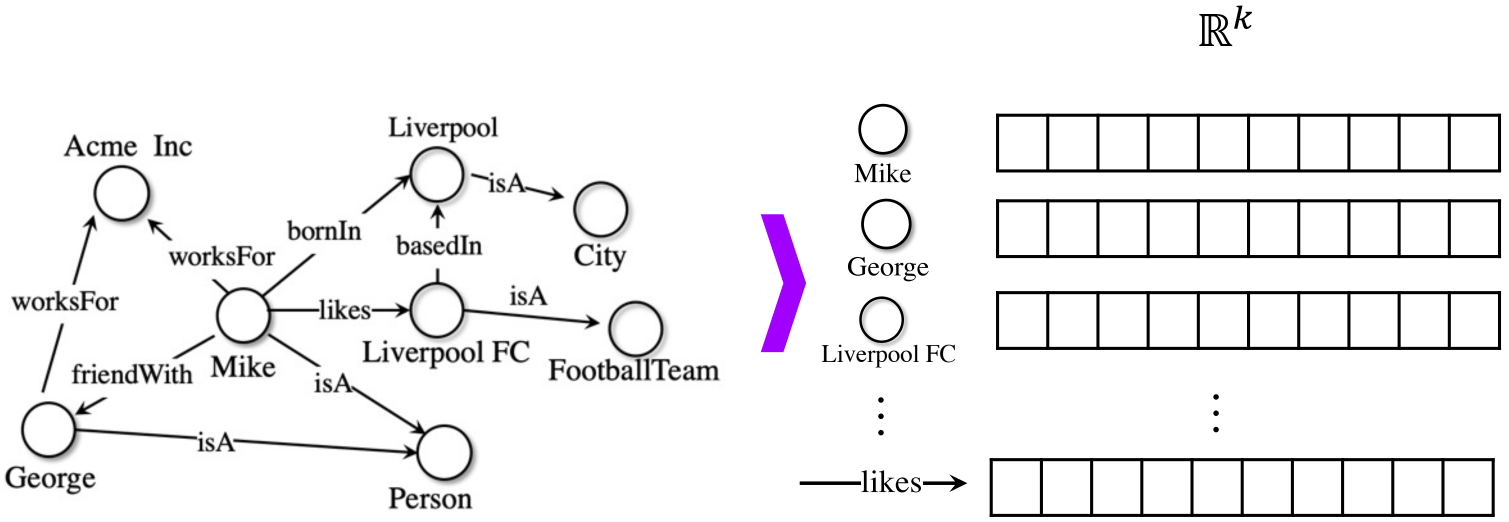

Downstream Task

Node Representation/Graph Feature based Methods

PRA, LINE, DeepWalk, node2vec

Graph Neural Networks (GNNs)

GCNs, Graph Attention Networks

Knowledge Graph Embeddings (KGE)

TransE, DistMult, ComplEx, ConvE

Scope of this tutorial

For a complete overview of graph feature-based models and GNNs:
[Hamilton & Sun 2019]
[Hamilton 2020]

Knowledge Graph Embeddings (KGE)

Automatic, supervised learning of **embeddings**, i.e. projections of entities and relations into a continuous low-dimensional space $\mathbb{R}^k$.

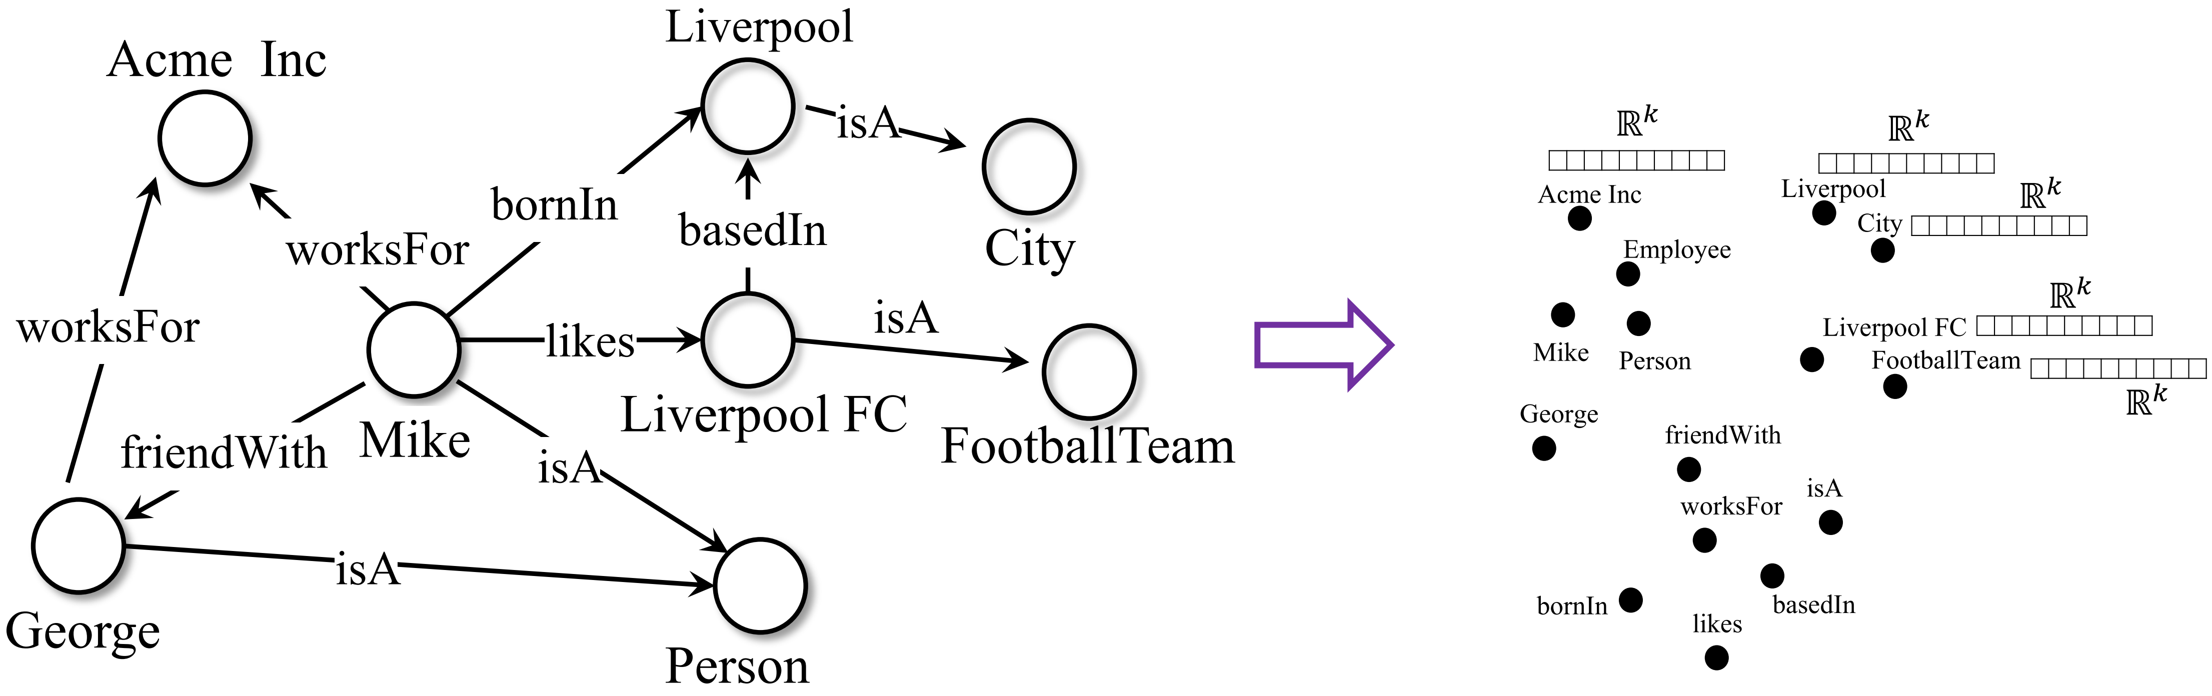

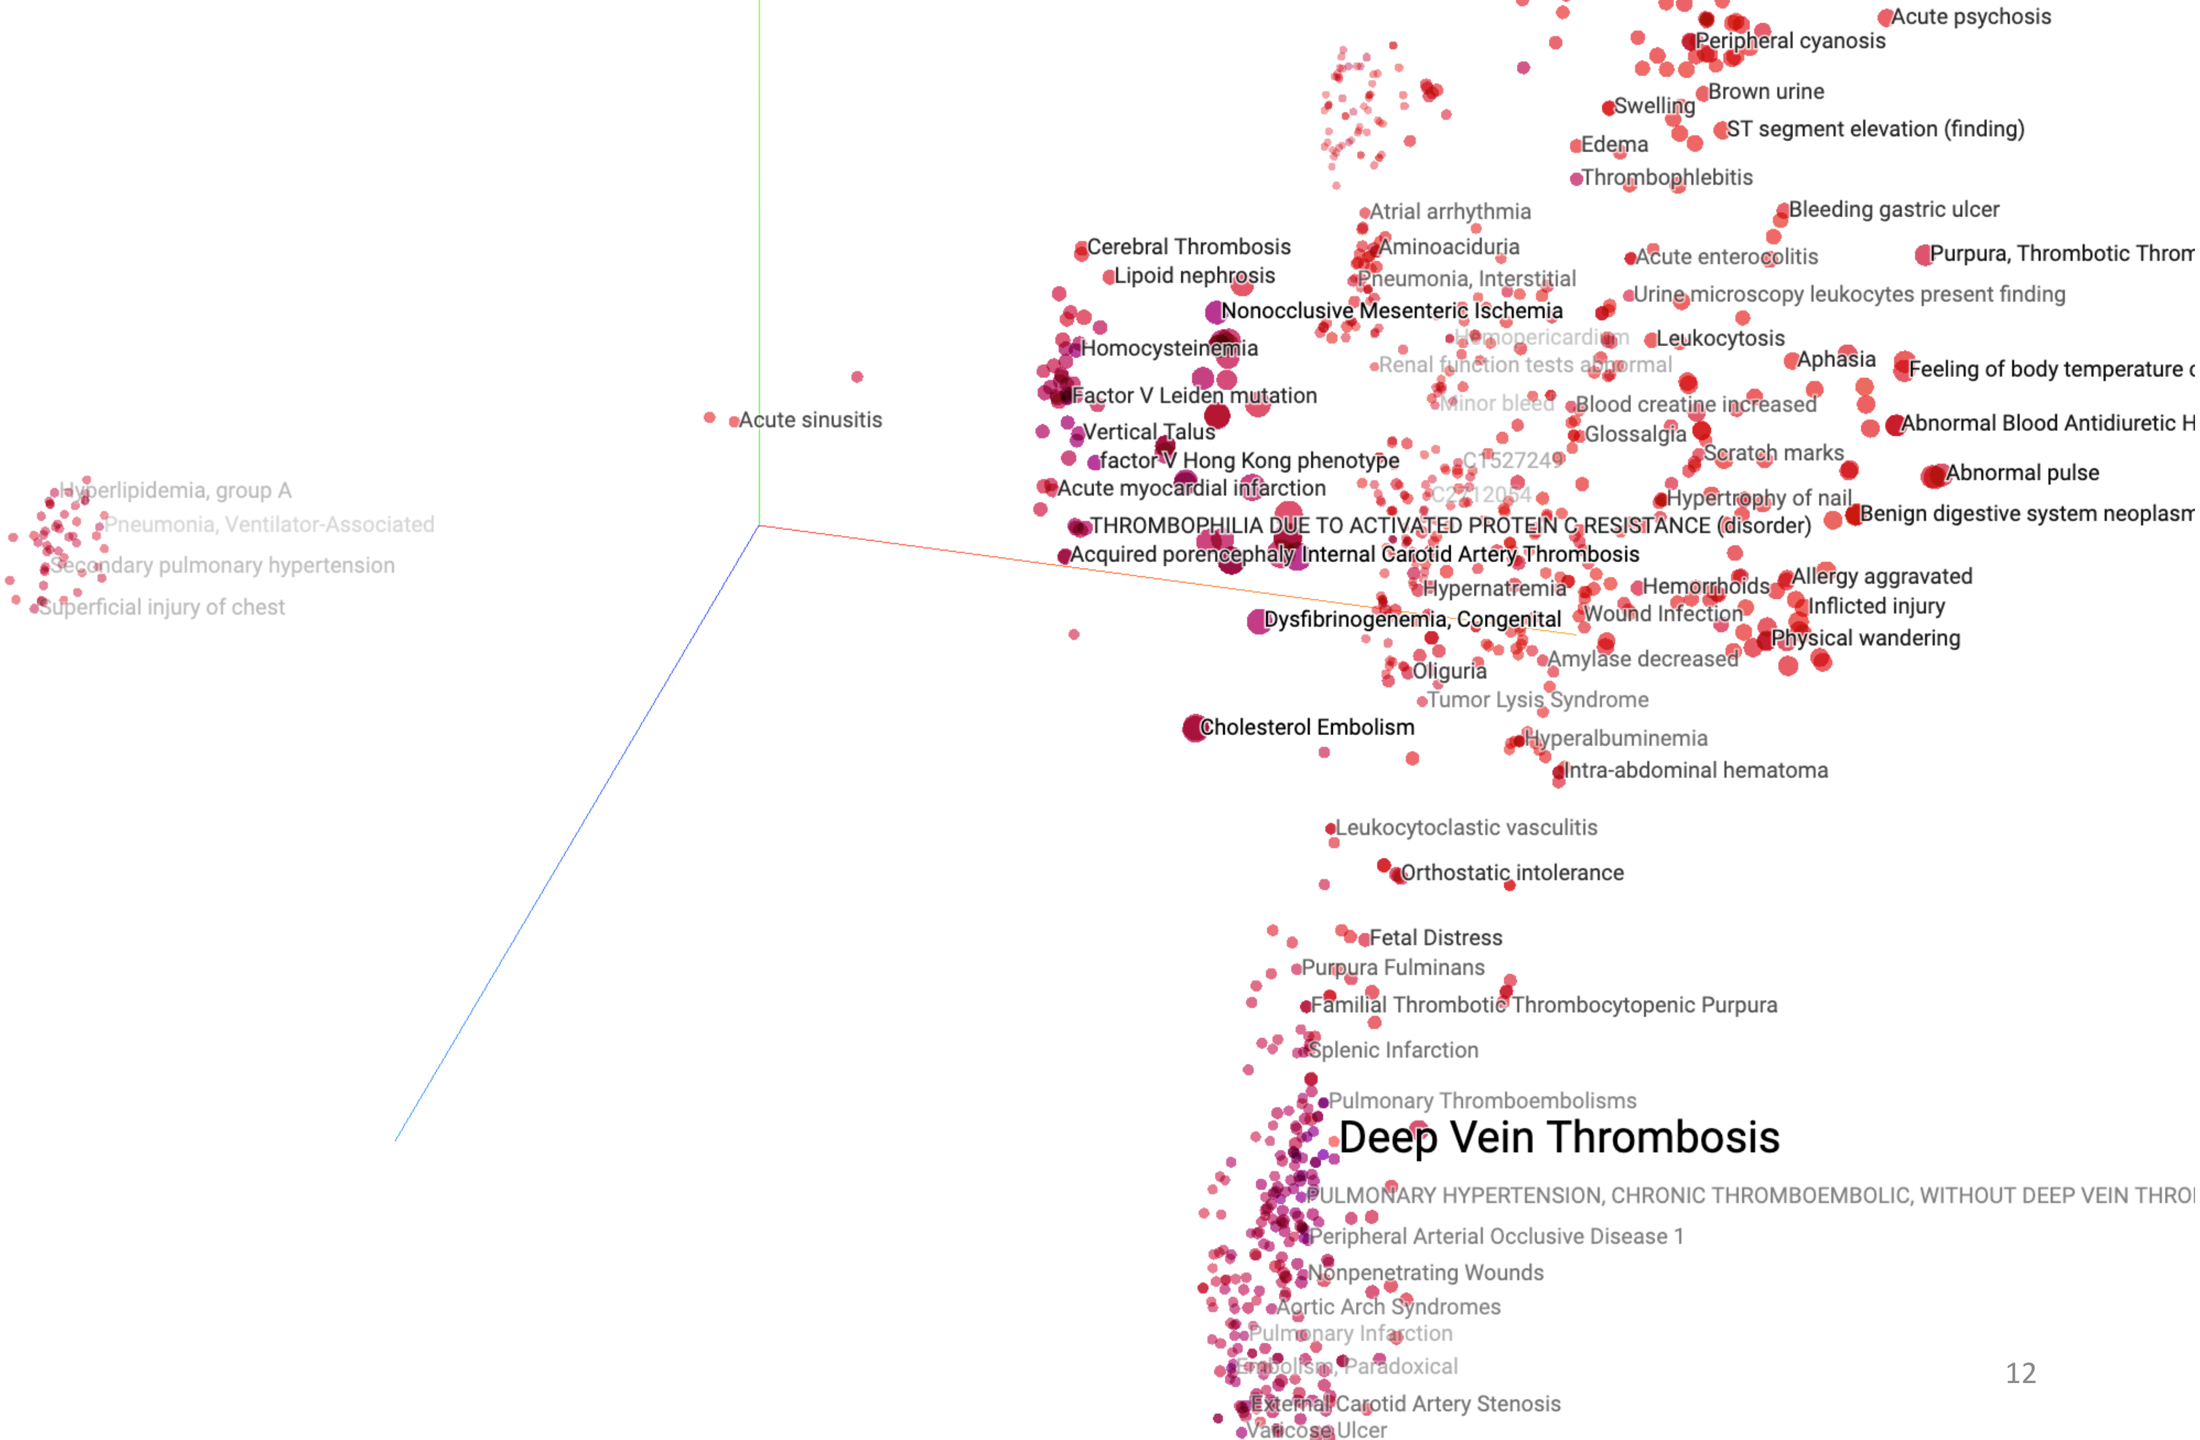
Acute psychosis
Peripheral cyanosis
Brown urine
Swelling
ST segment elevation (finding)
Edema
Thrombophlebitis
Atrial arrhythmia
Bleeding gastric ulcer
Cerebral Thrombosis
Aminoaciduria
Lipoid nephrosis
Pneumonia, Interstitial
Acute enterocolitis
Urine microscopy leukocytes present finding
Nonocclusive Mesenteric Ischemia
Hemopericardium
Leukocytosis
Homocysteinemia
Renal function tests abnormal
Aphasia
Feeling of body temperature c
Factor V Leiden mutation
Minor bleed
Blood creatine increased
Abnormal Blood Antidiuretic H
Acute sinusitis
Vertical Talus
Glossalgia
factor V Hong Kong phenotype
C1527249
Scratch marks
Abnormal pulse
Hyperlipidemia, group A
Acute myocardial infarction
C2712054
Hypertrophy of nail
Benign digestive system neoplasm
Pneumonia, Ventilator-Associated
THROMBOPHILIA DUE TO ACTIVATED PROTEIN C RESISTANCE (disorder)
Acquired porencephaly
Internal Carotid Artery Thrombosis
Secondary pulmonary hypertension
Hypernatremia
Hemorrhoids
Allergy aggravated
Superficial injury of chest
Inflicted injury
Wound Infection
Dysfibrinogenemia, Congenital
Physical wandering
Amylase decreased
Oliguria
Tumor Lysis Syndrome
Cholesterol Embolism
Hyperalbuminemia
Intra-abdominal hematoma
Leukocytoclastic vasculitis
Orthostatic intolerance
Fetal Distress
Purpura Fulminans
Familial Thrombotic Thrombocytopenic Purpura
Splenic Infarction
Pulmonary Thromboembolisms
Deep Vein Thrombosis
PULMONARY HYPERTENSION, CHRONIC THROMBOEMBOLIC, WITHOUT DEEP VEIN THRO
Peripheral Arterial Occlusive Disease 1
Nonpenetrating Wounds
Aortic Arch Syndromes
Pulmonary Infarction
Embolism, Paradoxical
External Carotid Artery Stenosis
Varicose Ulcer
Purpura, Thrombotic Throm

KGE Design Rationale: Capture KG Patterns

Symmetry <Alice marriedTo Bob>

Asymmetry <Alice childOf Jack>

Inversion
<Alice childOf Jack>
<Jack fatherOf Alice>

Composition
<Alice childOf Jack>
<Jack siblingOf Mary>
<Alice nieceOf Mary>

But also:
- Hierarchies
- Type constraints
- Transitivity
- Homophily
- Long-range dependencies

Popular KGE models in recent published literature

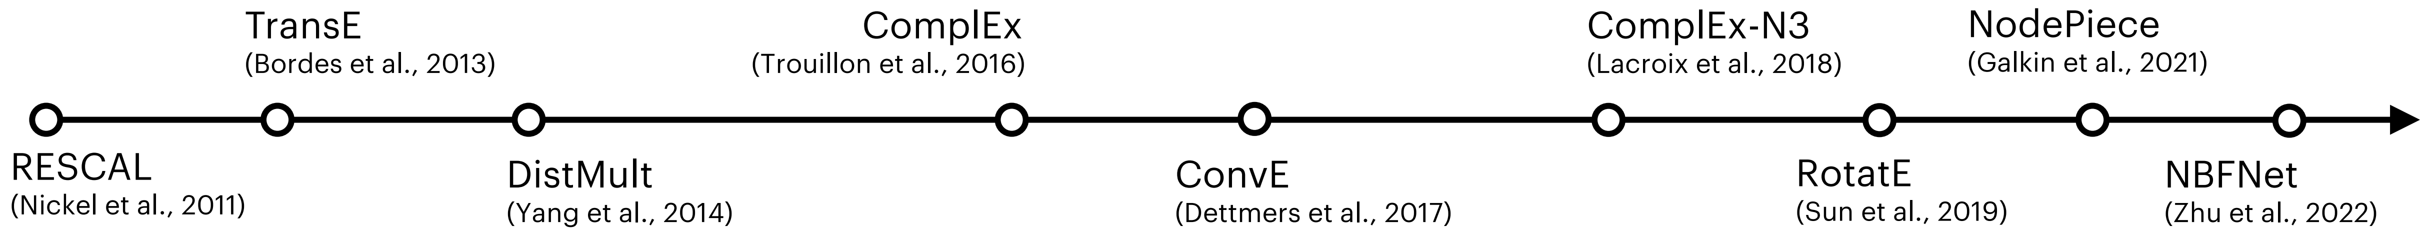

At a Glance

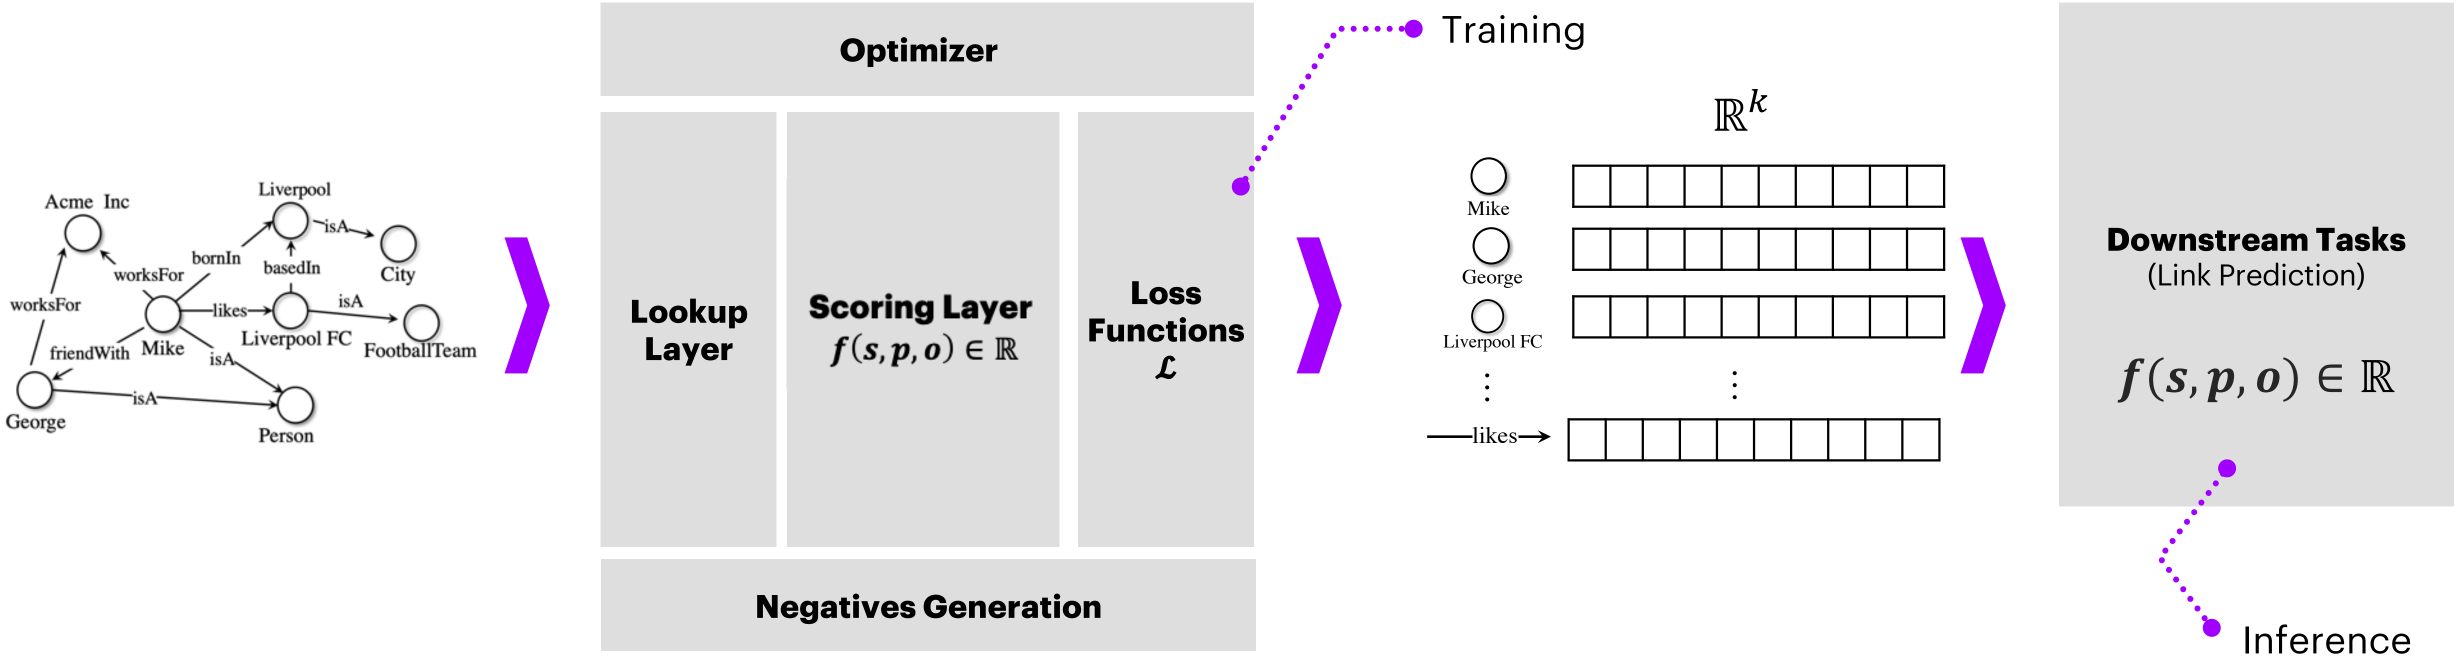
Acme Inc
Liverpool
City
worksFor
bornIn
basedIn
isA
worksFor
likes
isA
friendWith
Mike
Liverpool FC
FootballTeam
isA
George
isA
Person
Optimizer
Lookup Layer
Scoring Layer
$f(s,p,o) \in \mathbb{R}$
Loss Functions
$\mathcal{L}$
Negatives Generation
Training
$\mathbb{R}^k$
Mike
George
Liverpool FC
likes
Downstream Tasks
(Link Prediction)
$f(s,p,o) \in \mathbb{R}$
Inference

Anatomy of a Knowledge Graph Embedding Model

- Knowledge Graph (KG) $\mathcal{G}$
- Scoring function for a triple $f(t)$
- Loss function $\mathcal{L}$ (Translation-based, Factorization-based, Deep)
- Optimization algorithm
- Negatives generation strategy

Scoring function f

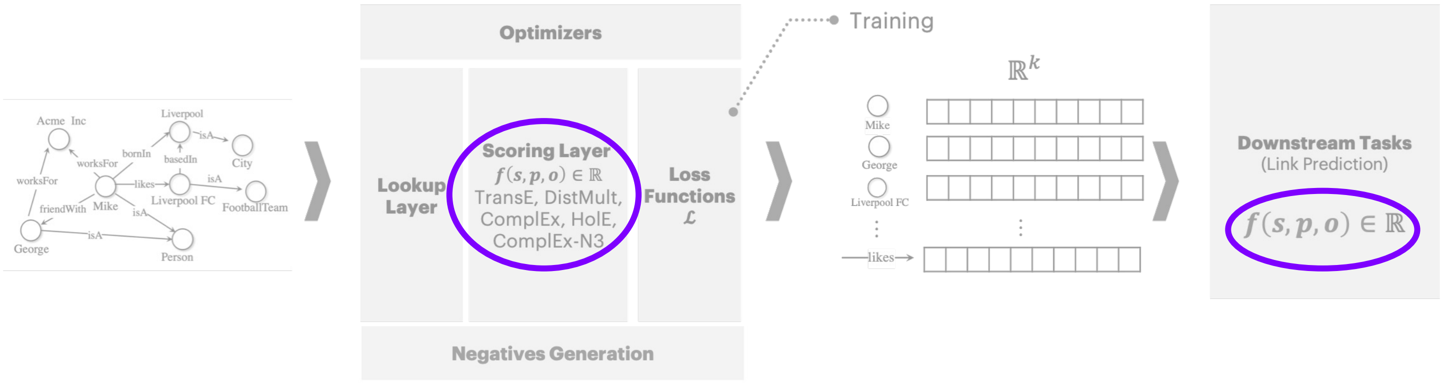

f assigns a score to a triple (s, p, o)

High score = triples is very likely to be factually correct

Translation-based Scoring Functions

- **TransE**: Translating Embeddings [Bordes et al. 2013]

$$f_{TransE} = -||(\mathbf{e}_s + \mathbf{r}_p) - \mathbf{e}_o||_n$$

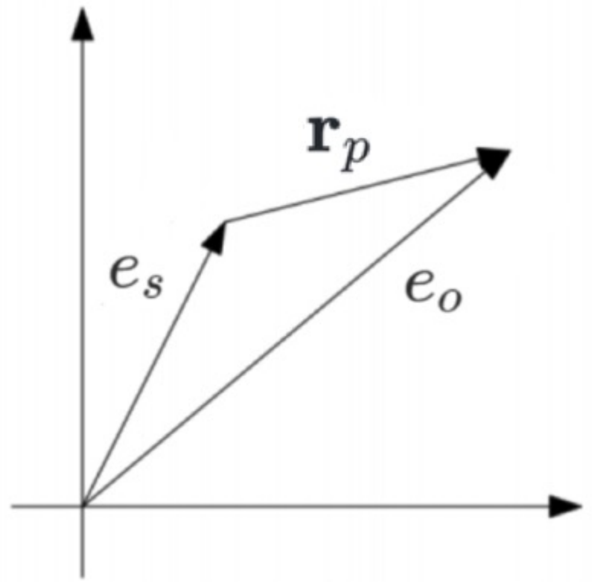

Translation-based Scoring Functions

- **RotatE**: relations modelled as *rotations* in complex space $\mathbb{C}$: element-wise product between complex embeddings. [Sun et al. 2019]

$$f_{RotatE} = -||\mathbf{e}_s \circ \mathbf{r}_p - \mathbf{e}_o||_n$$

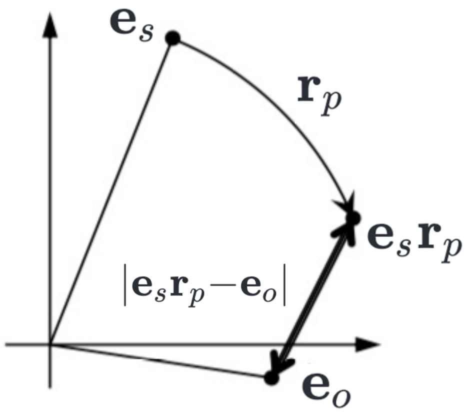

Factorization-based Scoring Functions

- **RESCAL**: low-rank factorization with tensor product

[Nickel et al. 2011]

$$f_{RESCAL} = \mathbf{e}_s^T \mathbf{W}_r \mathbf{e}_o$$

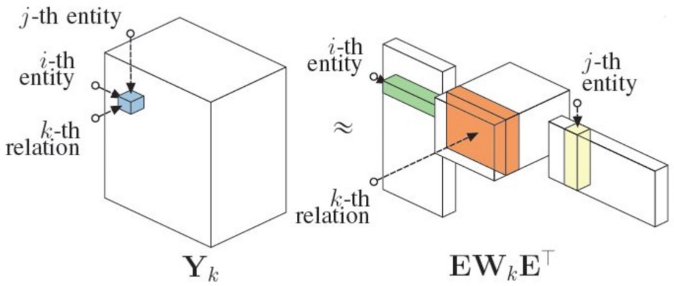

- **DistMult**: bilinear diagonal model. Dot product.

[Yang et al. 2015]

$$f_{DistMult} = \langle \mathbf{r}_p, \mathbf{e}_s, \mathbf{e}_o \rangle$$

- **ComplEx**: Complex Embeddings (Hermitian dot product):
  (i.e. extends DistMult with dot product in $\mathbb{C}$)

$$f_{ComplEx} = Re(\langle \mathbf{r}_p, \mathbf{e}_s, \overline{\mathbf{e}_o} \rangle)$$

[Trouillon et al. 2016]

"Deeper" Scoring Functions

- **ConvE**: reshaping + convolution

[Dettmers et al. 2017]

$$f_{ConvE} = \langle \sigma\left(\underbrace{vec}_{\text{2D reshaping}}\left(\overbrace{g}^{\text{Non-linearity}}\left(\left[\overline{\mathbf{e}_s};\overline{\mathbf{r}_p}\right] \underbrace{*}_{\text{Linear convolution}} \Omega\right)\right)\mathbf{W}\right)\mathbf{e}_o\rangle$$

- **ConvKB**: convolutions and dot product

[Nguyen et al. 2018]

$$f_{ConvKB} = concat\left(g\left(\left[\mathbf{e}_s,\mathbf{r}_p,\mathbf{e}_o\right]\right) * \Omega\right)\right) \cdot W$$

Computationally expensive!

Other Recent Models

- HolE [Nickel et al. 2016]
- SimplE [Kazemi et al. 2018]
- QuatE [Zhang et al. 2019]
- MurP [Balažević et al. 2019]
- NodePiece [Galkin et al. 2021]
- NBFNet [Zhú et al. 2022]
- …

Loss function $\mathcal{L}$

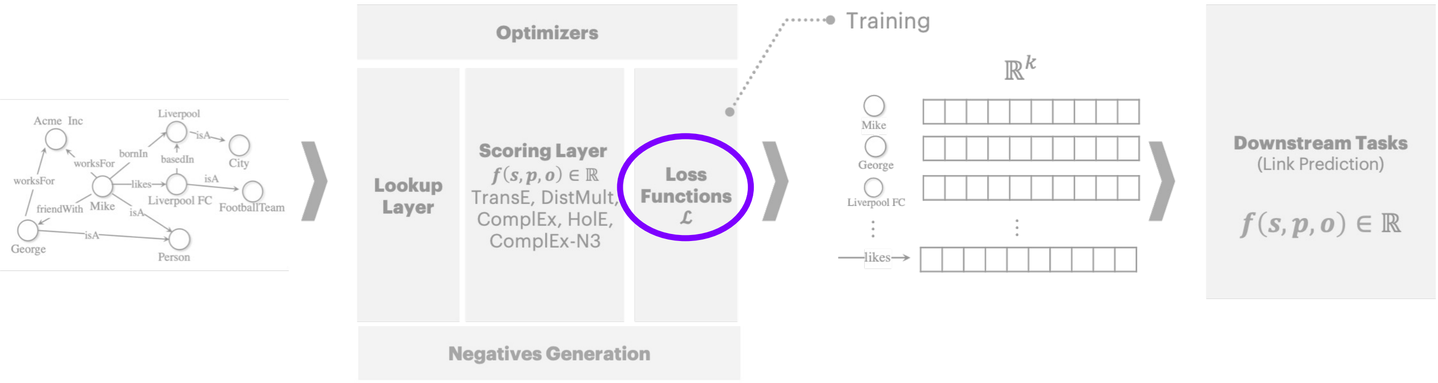

Pairwise Margin-Based Hinge Loss

Pays a penalty if score of positive triple < score of synthetic negative by a margin γ

$$\mathcal{L}(\Theta) = \sum_{t^+ \in \mathcal{G}} \sum_{t^- \in \mathcal{C}} max(0, [\gamma + \underbrace{f(t^-;\Theta)}_{\text{Score assigned to a \textbf{synthetic negative}}} - \underbrace{f(t^+;\Theta)}_{\text{Score assigned to \textbf{true} triple}}])$$

[Bordes et al. 2013]

Negatives Generation

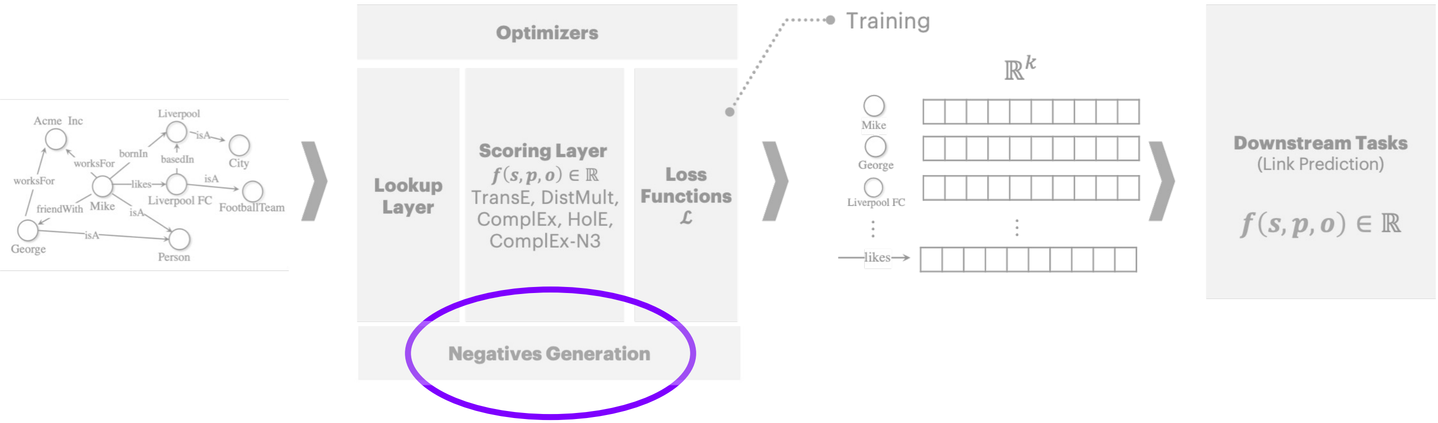

Where do negative examples come from? (i.e. false facts)

"Local Closed-World" Assumption: the KG is only *locally* complete

"Corrupted" versions of a triple as synthetic negatives:

$$\mathcal{C} = \{(\hat{s}, p, o) | \hat{s} \in \mathcal{E}\} \cup \{(s, p, \hat{o}) | \hat{o} \in \mathcal{E}\}$$

"corrupted subject"

"corrupted" object

The predicate is unaltered

Synthetic Negatives: Example

$$\mathcal{E} = \{Mike, Liverpool, AcmeInc, George, LiverpoolFC\}$$
$$\mathcal{R} = \{bornIn, friendWith\}$$

$t \in \mathcal{G} =$ `(Mike bornIn Liverpool)`

$\mathcal{C}_t =$

```
Mike      bornIn   AcmeInc
Mike      bornIn   LiverpoolFC
George    bornIn   Liverpool
AcmeInc   bornIn   Liverpool
```

Training Procedure and Optimizer

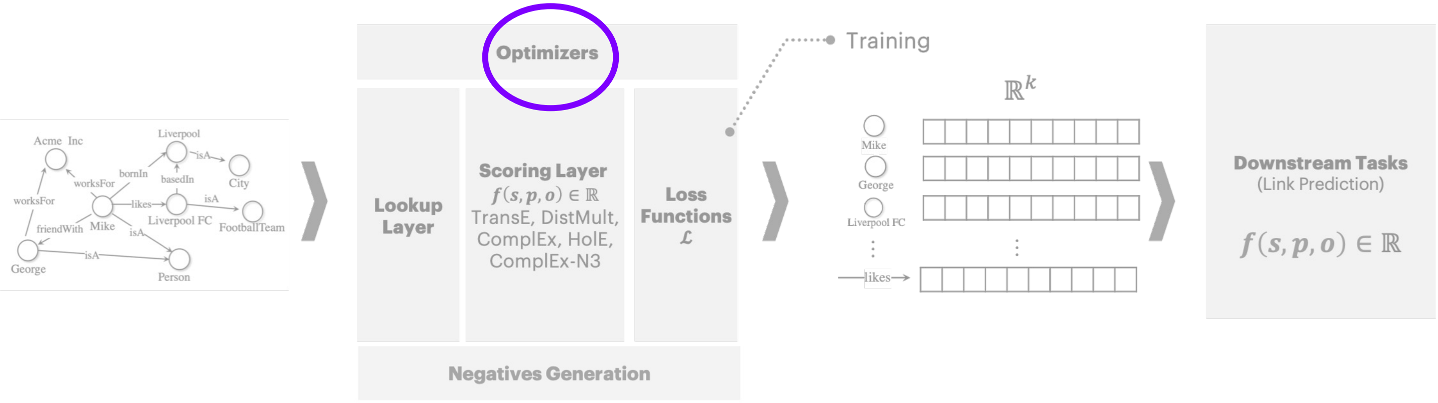

Optimizer: learn optimal parameters (e.g. embeddings). Off-the-shelf SGD variants: (AdaGrad, Adam)

$$\min_{\Theta} \mathcal{L}(\Theta)$$

Reciprocal Triples

Injection of reciprocal triples in training set.

`<Alice childOf Jack>`

`<Jack childOf`$^{-1}$ `Alice>`

[Dettmers et al. 2017]
[Lacroix et al. 2018]

Performance Evaluation

LINK PREDICTION / **TRIPLE CLASSIFICATION**

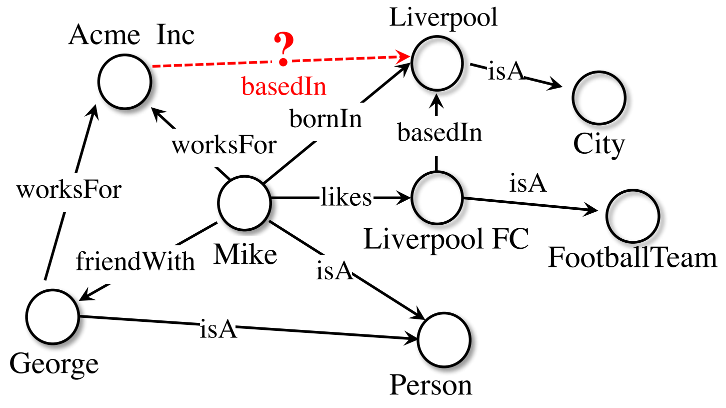

Assigning a score proportional to the likelihood that an unseen triple is true.

Link Prediction

- Learning to rank problem
- Information retrieval metrics
- No ground truth negatives in test set required

Triple Classification

- Binary classification task
- Binary classification metrics
- Test set requires positives and ground truth negatives

Same procedure used in training

Learning–To-Rank problem:

How well are positive triples ranked against **synthetic negatives** built under the **Local Closed World Assumption.**

Evaluation Metrics

Mean Rank (MR)

$$MR = \frac{1}{|Q|} \sum_{i=1}^{|Q|} rank_{(s,p,o)_i}$$

Mean Reciprocal Rank (MRR)

$$MRR = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{1}{rank_{(s,p,o)_i}}$$

Hits@N

$$Hits@N = \sum_{i=1}^{|Q|} 1 \text{ if } rank_{(s,p,o)_i} \leq N$$

Example: How unseen, test positive triples rank against **synthetic negatives**? (four negatives/positive)

s	p	o	score	rank	
Mike	friend_with	George	0.901	1	*
Mike	friend_with	Jim	0.345	2	
Acme	friend_with	George	0.293	3	
Mike	friend_with	Liverpool	0.201	4	
France	friend_with	George	0.156	5	

s	p	o	score	rank	
Mike	born_in	Leeds	0.789	1	
Mike	born_in	Liverpool	0.753	2	*
Mike	born_in	Germany	0.695	3	
George	born_in	Liverpool	0.456	4	
Mike	born_in	George	0.234	5	

Positive triples from test set

Test set = {
<Mike friend_with George>
<Mike born_in Liverpool>
}

$$MR = 1.5$$
$$MRR = 0.75$$
$$Hits@1 = 0.5$$
$$Hits@3 = 1.0$$

TRUSTING PREDICTIONS: CALIBRATION

Probabilities Generated by off-the-shelf KGE models are uncalibrated!

- **Mistrust** in model discoveries
- **Poor Interpretability** in high-stakes scenarios (i.e. drug-target discovery)

Can we calibrate KGE models? Yes, and that leads to more **trustworthy and interpretable** predictions.

Probabilities returned by off the shelf models do not match the actual fraction of positives!

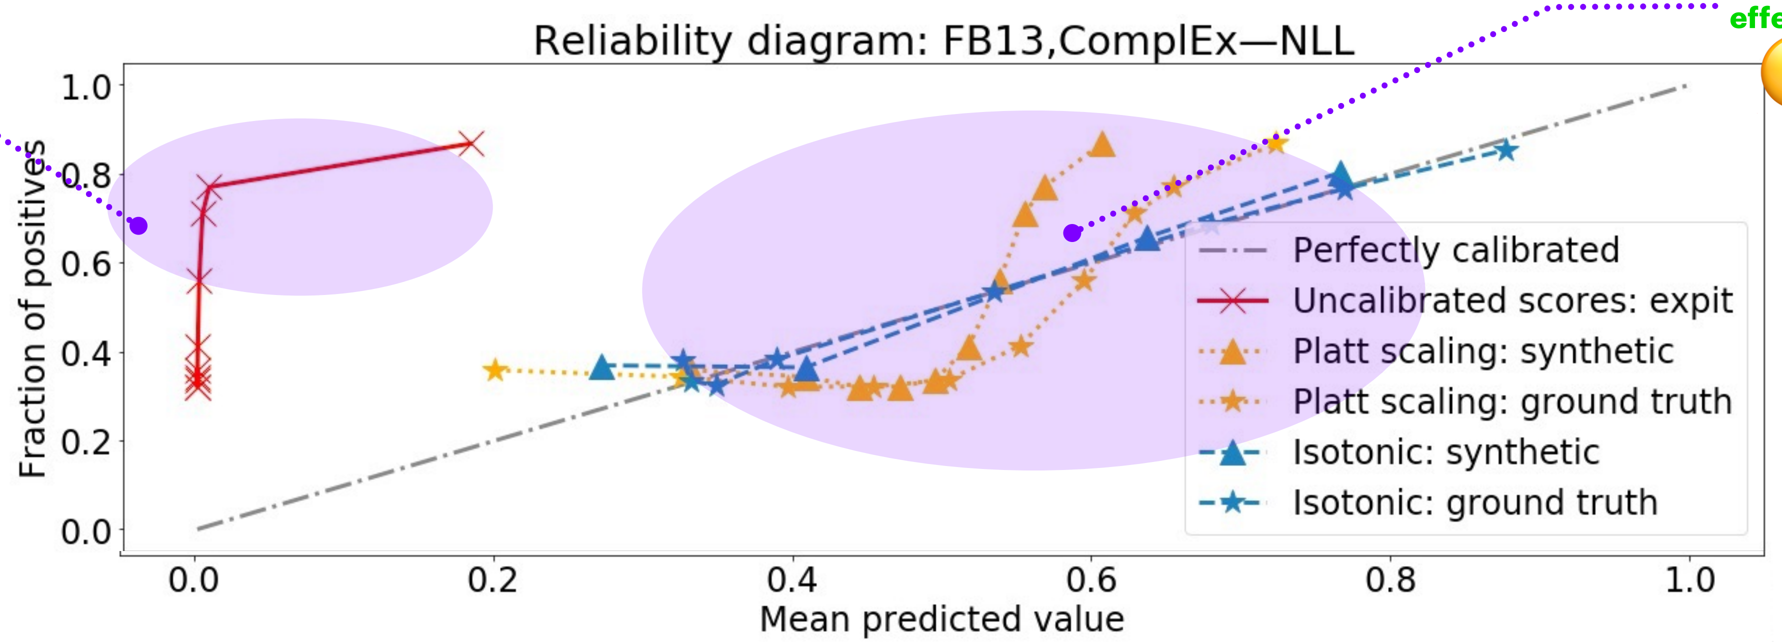

Calibration is effective!

[Tabacof & Costabello ICLR 2020]

arXiv:1912.10000

EXPLAINING PREDICTIONS

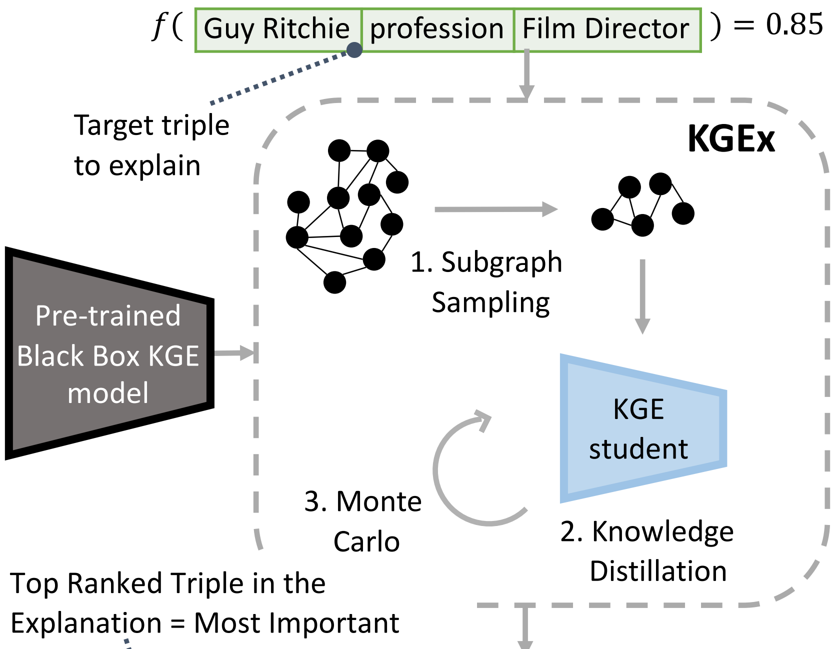

Explanation

1	Guy Ritchie	directed	Sherlock Holmes
2	Guy Ritchie	married	Madonna
3	Madonna	profession	Film Producer
		⋮	

[Baltatzis & Costabello LoG-2023]

arXiv:2310.01065

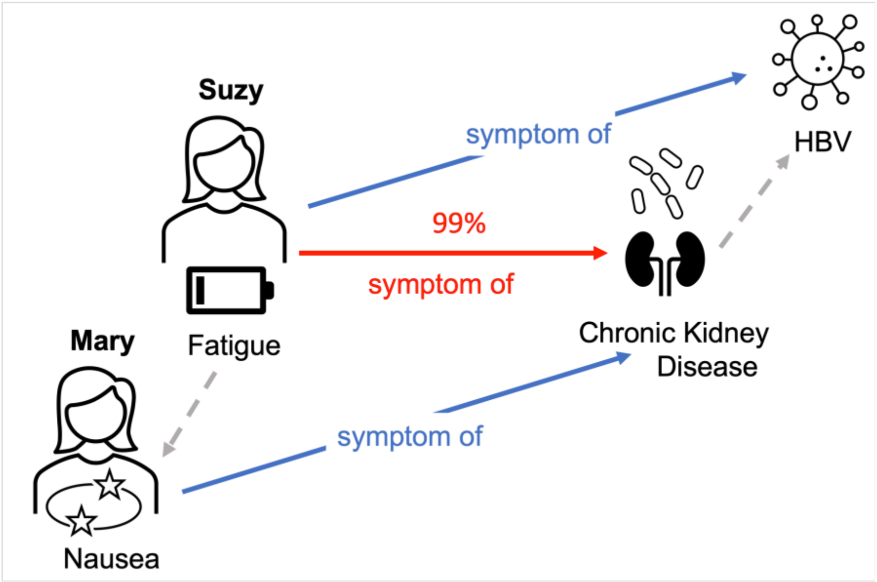

Figure 1: To support prediction of the target statement we identify influential examples by probing the knowledge base constrained w.r.t. the latent-space. This example is drawn from the Fb15k-237 dataset. Predicted plausability score was 99%, and two most influential examples were retrieved as an explanation with the following ranks: 1st: $Nausea \rightarrow symptomOf \rightarrow ChronicKidneyDisease$, 2nd: $Fatigue \rightarrow symptomOf \rightarrow HBV$

[Janik & Costabello LoG-2023]

arXiv:2212.02651

MULTIMODAL KNOWLEDGE GRAPHS

Many real-world graphs include **multi-modal attributes**.

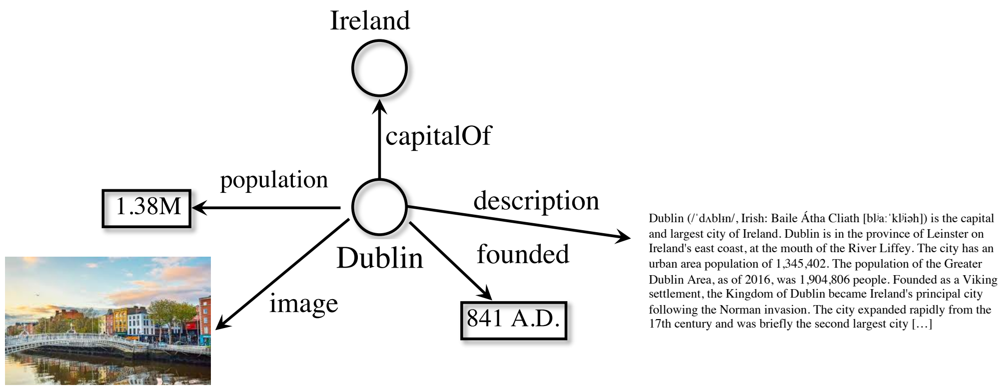

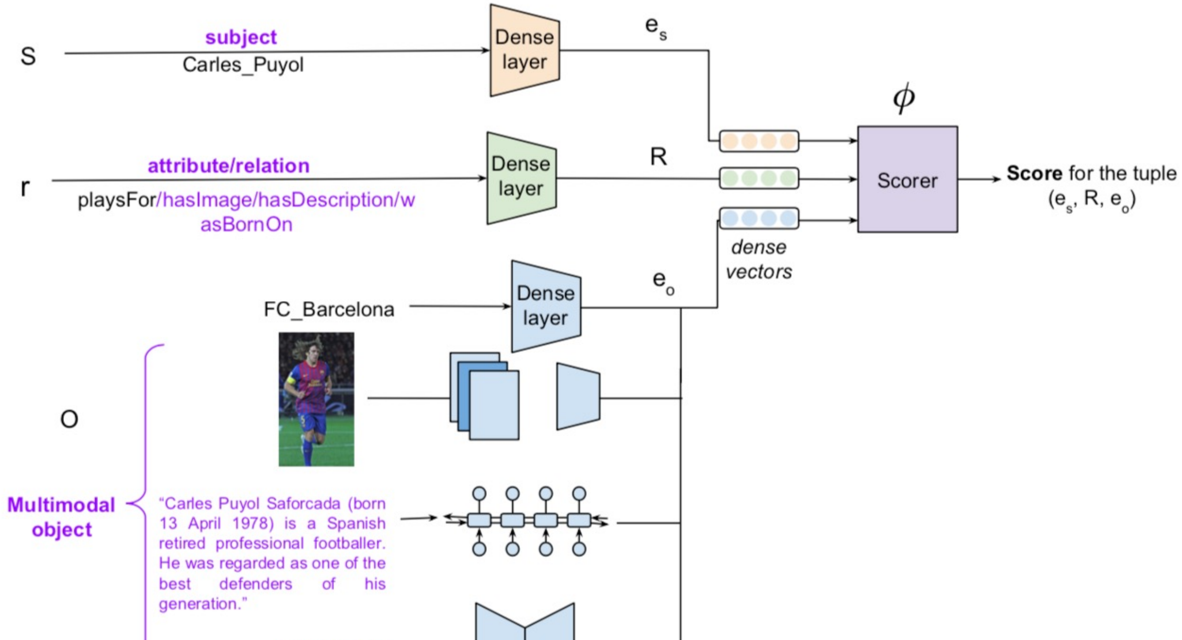

[Pezeshkpour et al. 2018]

[Gesese et al. 2019] surveys recent literature

MULTIMODAL KNOWLEDGE GRAPHS

Many real-world graphs includes numeric informa on edges (e.g. “strength”)

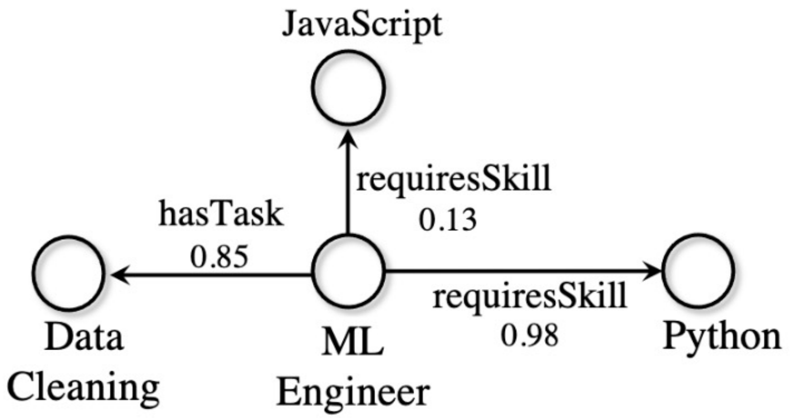

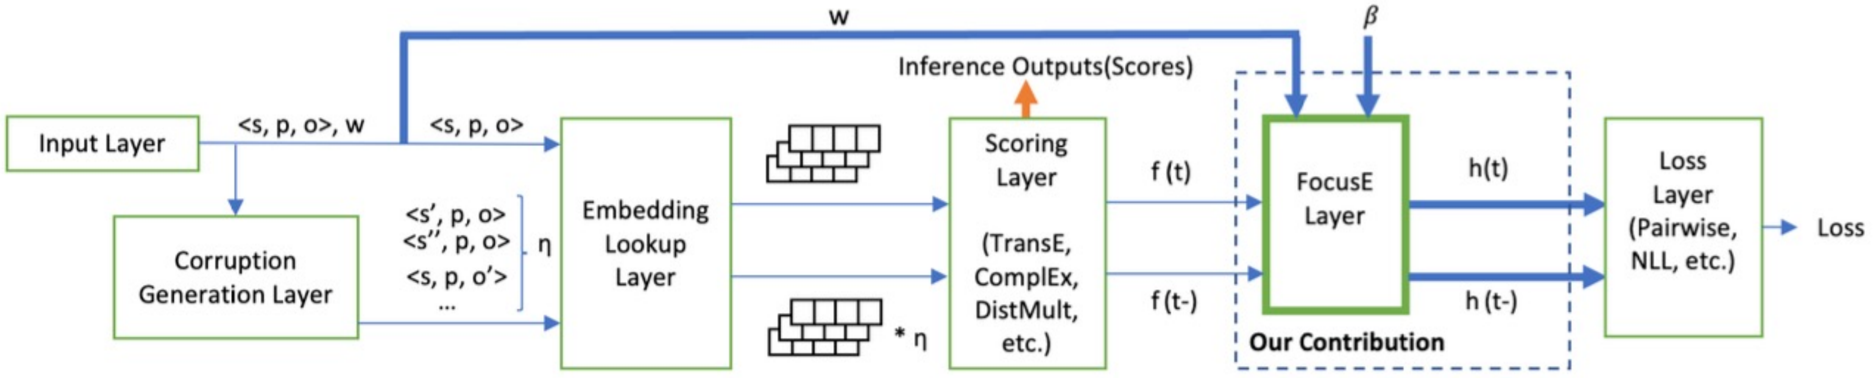

[Pai & Costabello IJCAI-21]

arXiv:2105.08683

TEMPORAL KNOWLEDGE GRAPHS

Many real-world graphs represents timestamped concepts.

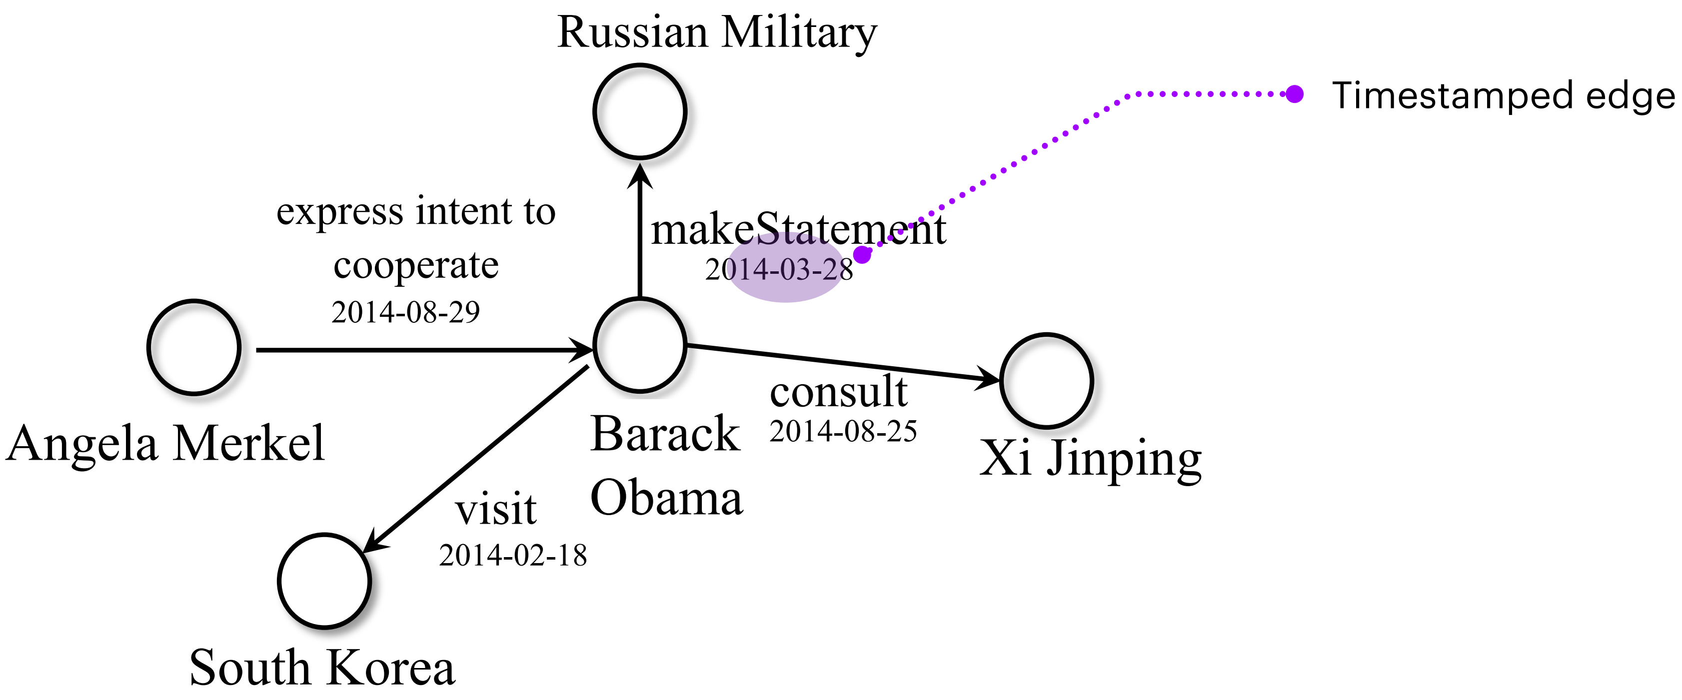

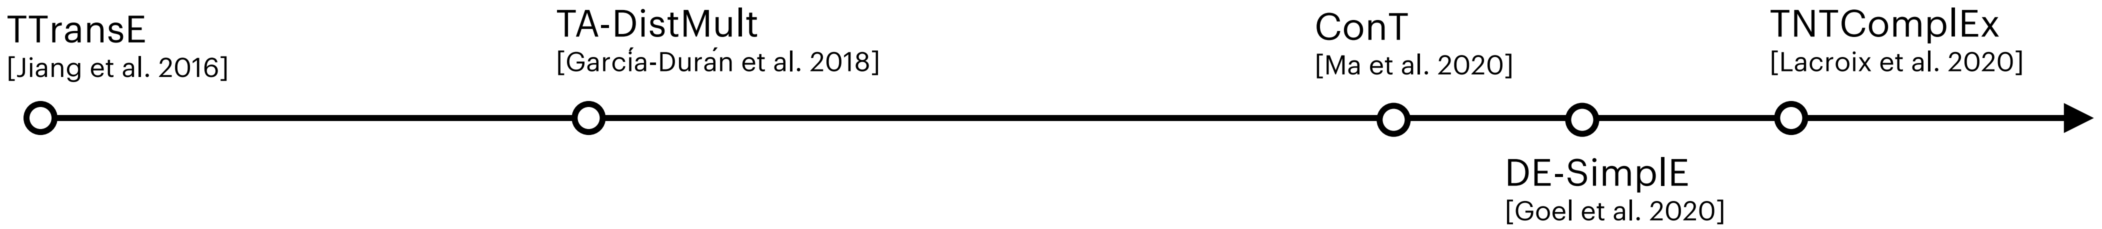

ROBUSTNESS TO ADVERSARIAL ATTACKS

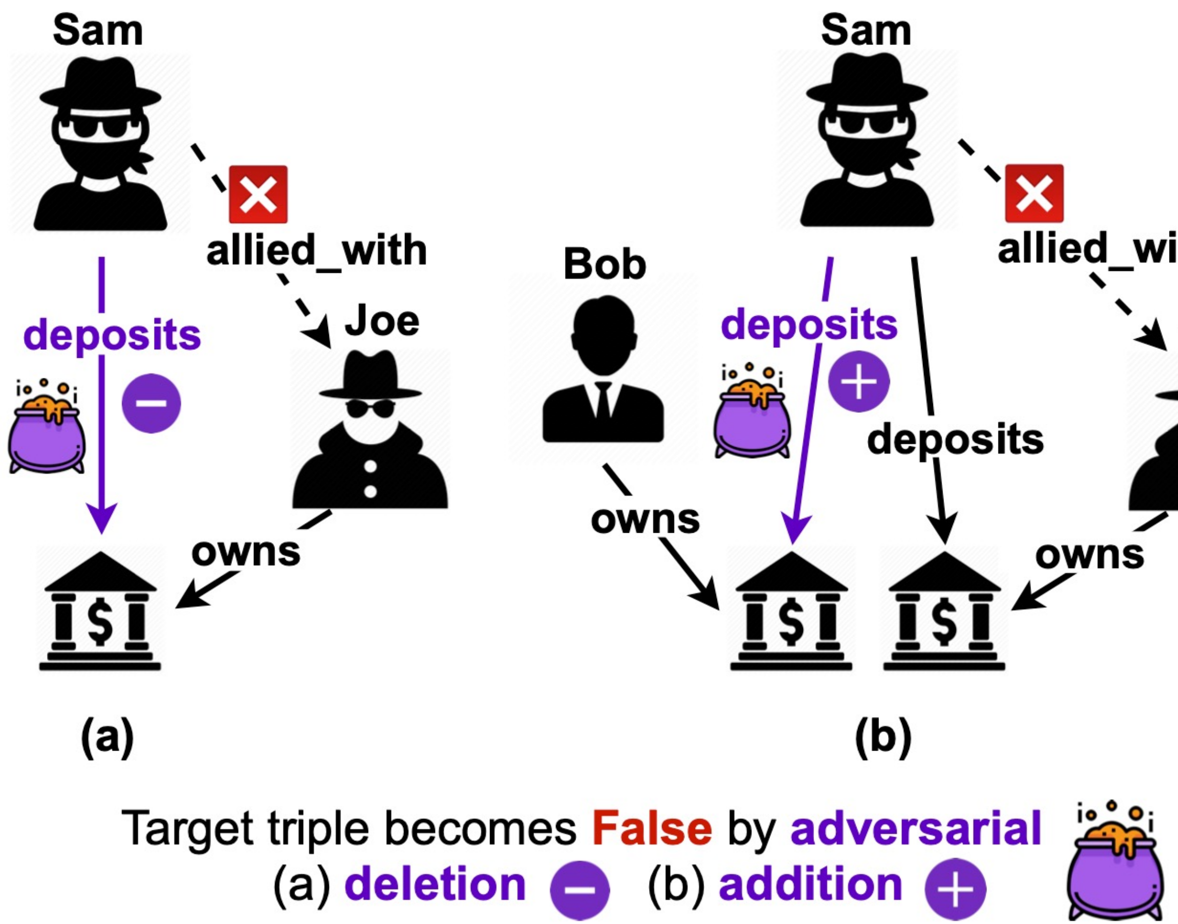

[Bhardwaj EMNLP-21]
[Bhardwaj ACL-21]

NEURAL GRAPH DATABASES: MULTI HOP QUESTION ANSWERING

"Which directors directed actors that won either an Oscar or an Emmy?"

$$?D : \exists A \,.\, \mathrm{directs}(D, A) \wedge \left[\mathrm{prize}(A, \mathrm{Oscar}) \vee \mathrm{prize}(A, \mathrm{Emmy})\right]$$

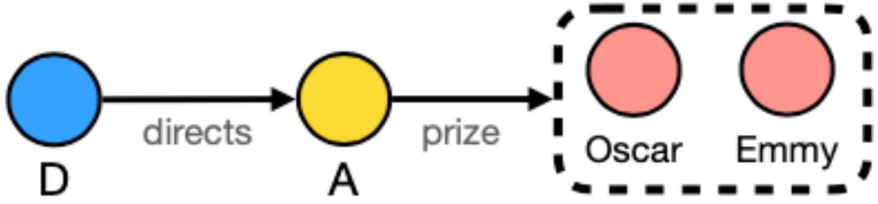

[Arakelyan et al ICLR 2021]

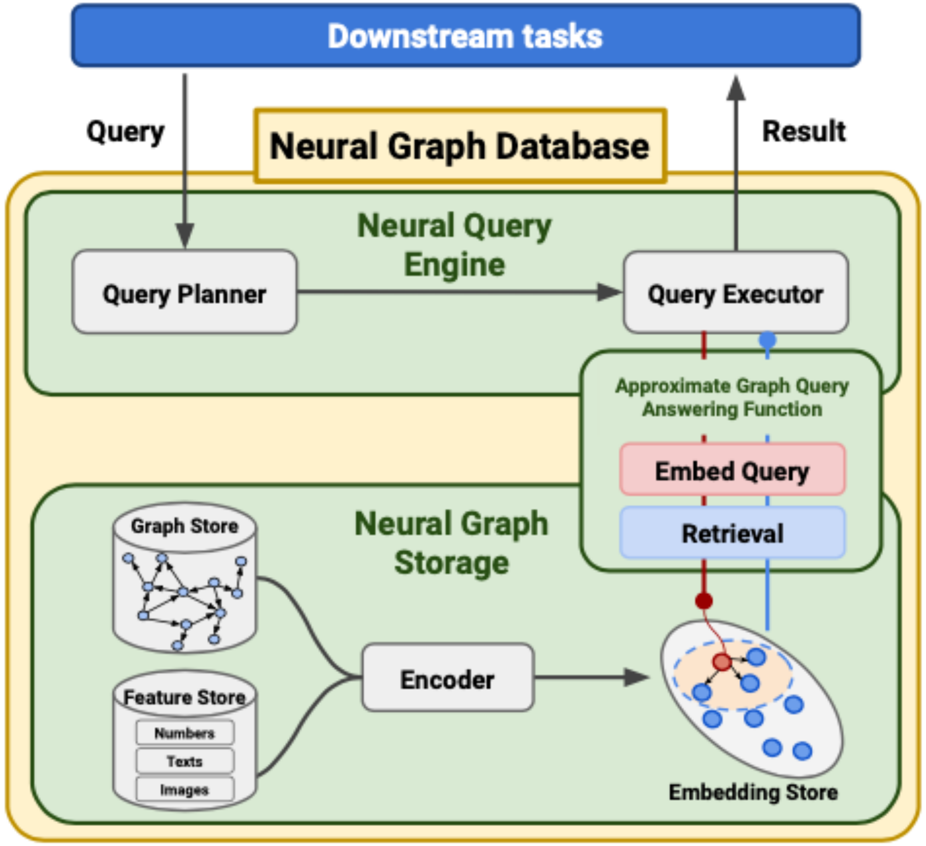

LLM-KGE Interplay: Joint embeddings

End-to-end architectures that “fuse” text embeddings with graph embeddings to increase predictive power

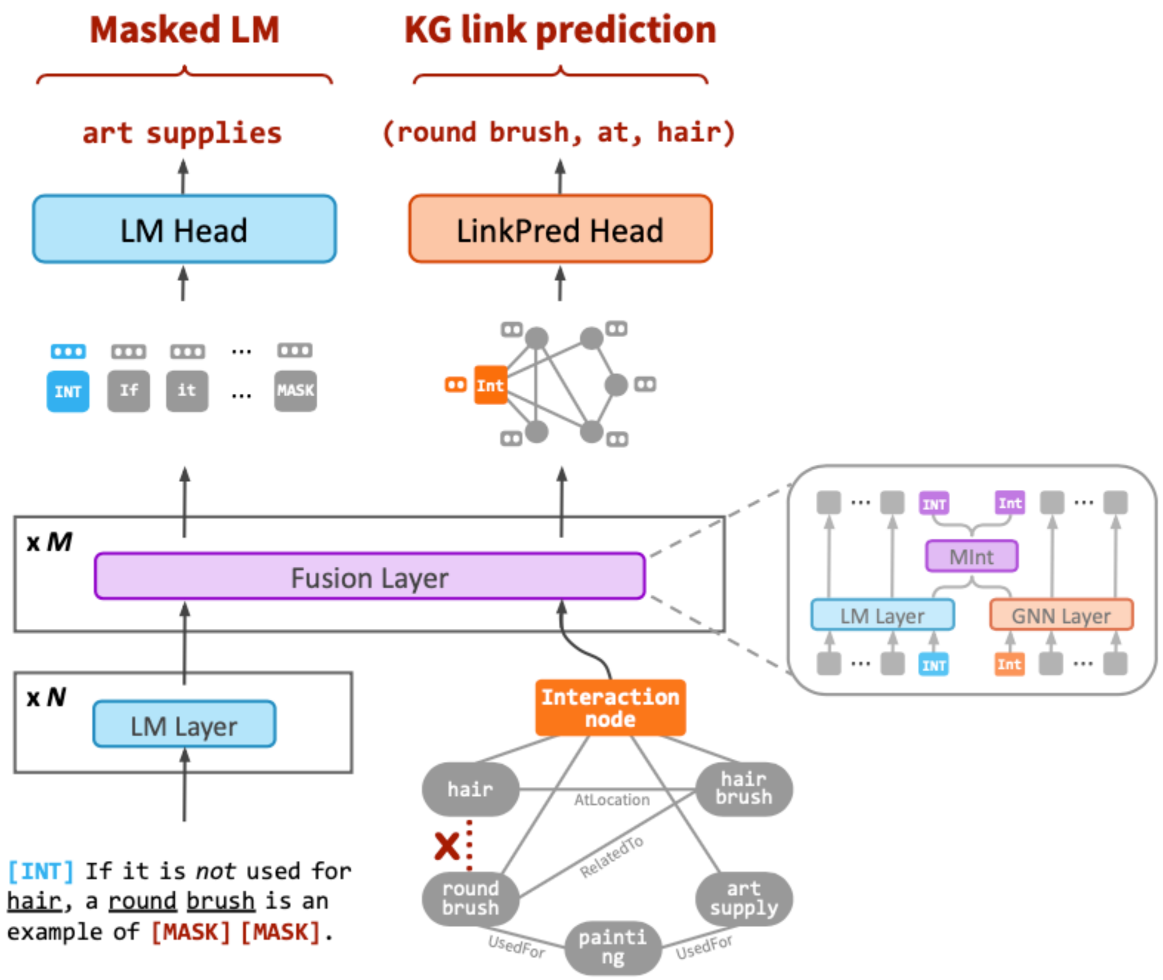

[Yasunaga et al 2022]

AmpliGraph

ampligraph.org

`pip install ampligraph`

OPEN SOURCE PYTHON LIBRARY FOR GRAPH REPRESENTATION LEARNING WITH KNOWLEDGE GRAPH EMBEDDINGS

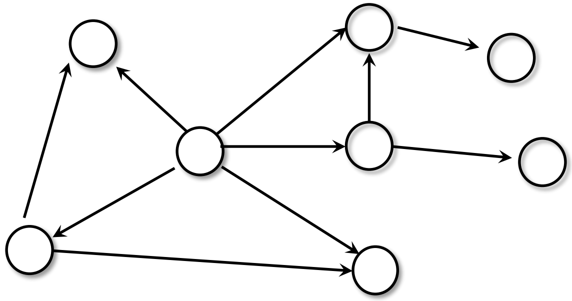

AmpliGraph

Link Prediction

Link-based Clustering

Collective Entity Matching

INDUSTRIAL APPLICATIONS AT ACCENTURE LABS BIOINNOVATION

Pharma
Drug-Target Interaction Discovery

Oncology
Early Lung cancer patients relapse prediction

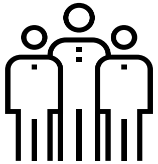

DRUG DEVELOPMENT

Pharma
Drug-Target Interaction Discovery

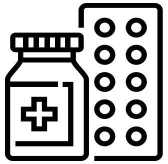

Target Identification → Target Validation → Pre Clinical Trials → Human Clinical Trials → Regulatory Approvals and Manufacturing

7 - 15 years

Drug Targets 101

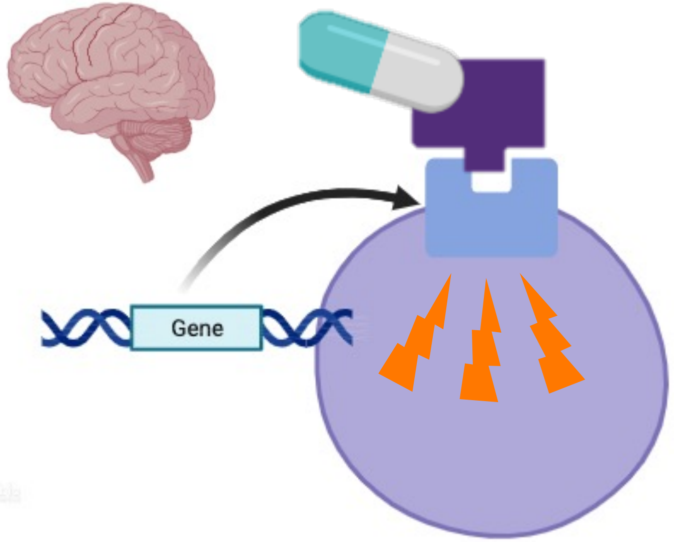

How to find genes that express promising targets?

Problem: 21k+ Genes to Choose From!

ADRA2A
CACNA1F
GRIA4
HRH1
GABRA2
ADORA2B
CACNA1I
GABRG3
ADRA2C
ADORA1
.
.
.

Knowledge Graphs and GraphML To The Rescue

We Use KGEs to Rank Order Genes by Importance

1 HRH1
2 ADRA2A
3 CACNA1F
4 ADORA2B
5 GABRG3
6 ADORA1
7 GRIA4
8 CACNA1I
9 GABRA2
10 ADRA2C

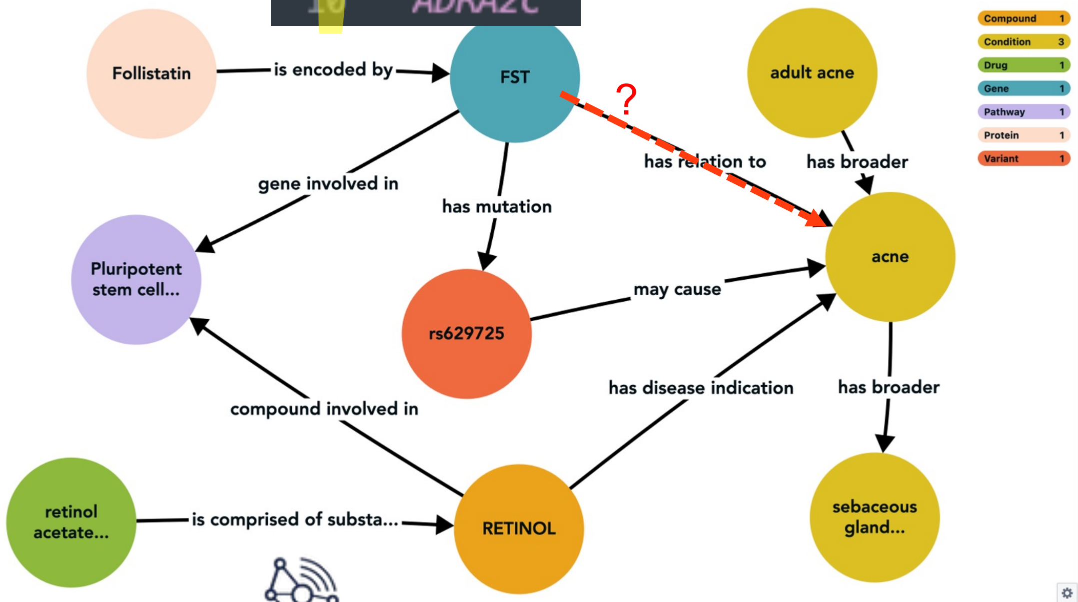

AmpliGraph

Impact

Success rate of 77% in identifying known gene-headache associations.

80% faster than traditional methods

Pharma
Drug-Target Interaction Discovery

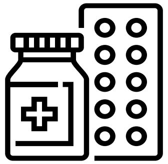

https://www.clarify2020.eu/

CLARIFY: Cancer Long Survivor Artificial Intelligence Follow-up

Oncology
Lung cancer patients relapse prediction

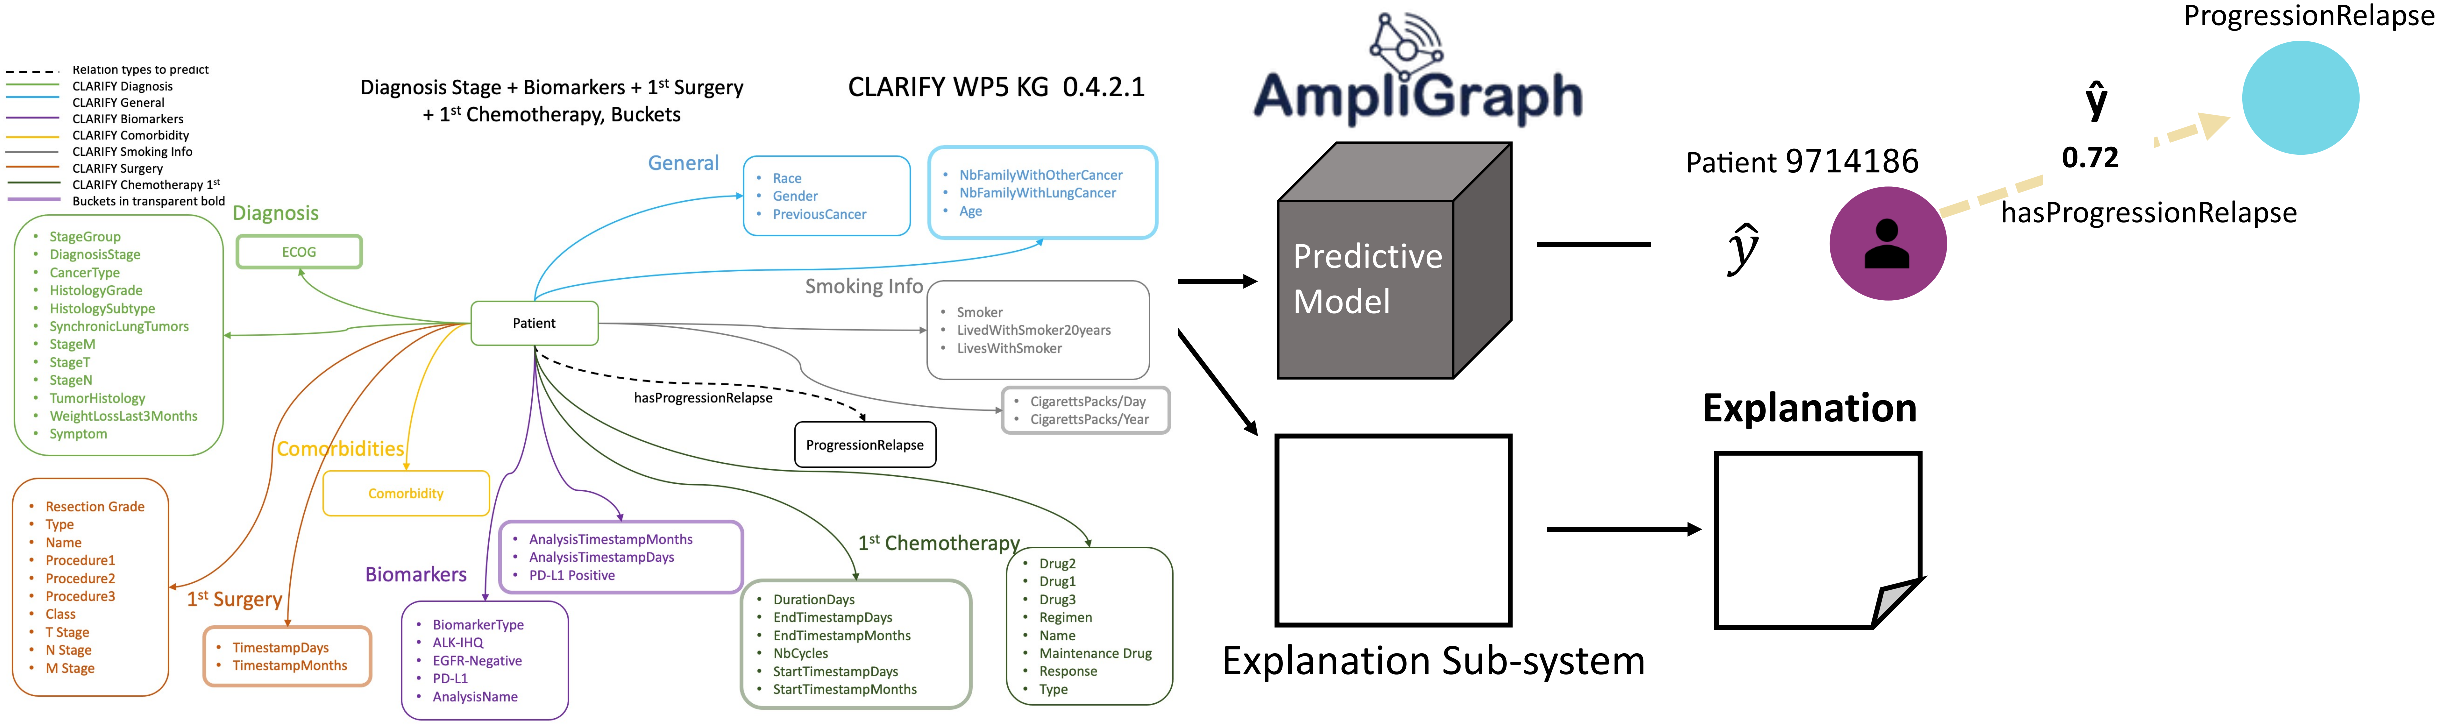

Accenture Labs Dublin

Luca Costabello
Research Scientist
Accenture Labs

luca.costabello@accenture.com
@lukostaz

github.com/Accenture/AmpliGraph

Q&A

REFERENCES

[Arakelyan 2021] Erik Arakelyan, Daniel Daza, Pasquale Minervini, Michael Cochez: Complex query Answering with Neural Link Predictors, ICLR 2021.

[Auer et al. 2007] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. Dbpedia: a nucleus for a web of open data. In The semantic web, 722–735. Springer, 2007.

[Bhardwaj ACL 2021] Peru Bhardwaj, John Kelleher, Luca Costabello and Declan O'Sullivan: Poisoning Knowledge Graph Embeddings via Relation Inference Patterns. ACL, 2021.

[Bhardwaj EMNLP 2021] Peru Bhardwaj, John Kelleher, Luca Costabello and Declan O'Sullivan: Adversarial Attacks on Knowledge Graph Embeddings via Instance Attribution Methods. EMNLP, 2021

[Bianchi et al. 2020] Federico Bianchi, Gaetano Rossiello, Luca Costabello, Matteo Palmonari, and Pasquale Minervini. "Knowledge Graph Embeddings and Explainable AI." arXiv preprint arXiv:2004.14843 (2020).

[Bollacker et al. 2008] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. Freebase: a collaboratively created graph database for structuring human knowledge. In Proceedings of the 2008 ACM SIGMOD international conference on Management of data, 1247–1250. AcM, 2008.

[Bordes et al. 2013] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. NIPS, 2013.

[Chen et al. 2019] Chen, Xuelu, Muhao Chen, Weijia Shi, Yizhou Sun, and Carlo Zaniolo. "Embedding uncertain knowledge graphs." In Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, pp. 3363-3370. 2019.

[Danqi et al. 2015] Danqi and Chen. Observed versus latent features for knowledge base and text inference. In 3rd Workshop on Continuous Vector Space Models and Their Compositionality. ACL, 2015.

[Dettmers et al. 2017] Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Convolutional 2d knowledge graph embeddings. AAAI, 2018

[García-Durán et al. 2017] Garcia-Duran, Alberto, and Mathias Niepert. "Kblrn: End-to-end learning of knowledge base representations with latent, relational, and numerical features." arXiv preprint arXiv:1709.04676 (2017).

[García-Durán et al. 2018] Alberto García-Durán, Sebastijan Dumancˇic´, and Mathias Niepert. Learning sequence encoders for temporal knowledge graph completion. arXiv preprint arXiv:1809.03202, 2018.

[Gesese et al. 2019] Gesese, Genet Asefa, Russa Biswas, Mehwish Alam, and Harald Sack. "A Survey on Knowledge Graph Embeddings with Literals: Which model links better Literal-ly?." arXiv preprint arXiv:1910.12507, 2019.

[Glorot & Bengio 2010] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In Proceedings of the thirteenth international conference on artificial intelligence and statistics, 249–256. 2010.

[Goodfellow et al. 2016] Goodfellow, Ian, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. Deep learning. Vol. 1. Cambridge: MIT press, 2016.

[Hamilton & Tang, 2019] William Hamilton, Jian Tang. Tutorial on Graph Representation Learning, AAAI, 2019.

[Hamilton 2020] William L. Hamilton. (2020). Graph Representation Learning. Morgan & Claypool, forthcoming.

[Hogan et al. 2020] Hogan, Aidan, Eva Blomqvist, Michael Cochez, Claudia d'Amato, Gerard de Melo, Claudio Gutierrez, José Emilio Labra Gayo et al. "Knowledge graphs." arXiv preprint arXiv:2003.02320 (2020).

[Kadlec 2017] Rudolf Kadlec, Ondrej Bajgar, and Jan Kleindienst. Knowledge base completion: baselines strike back. CoRR, 2017.

[Kristiadi et al. 2018] Kristiadi, Agustinus, Mohammad Asif Khan, Denis Lukovnikov, Jens Lehmann, and Asja Fischer. "Incorporating literals into knowledge graph embeddings." arXiv preprint arXiv:1802.00934 (2018).

[Lacroix et al. 2018] Timothee Lacroix, Nicolas Usunier, and Guillaume Obozinski. Canonical tensor decomposition for knowledge base completion. ICML, 2018.

[Li et al. 2018] Lisha Li and Kevin Jamieson. Hyperband: a novel bandit-based approach to hyperparameter optimization. Journal of Machine Learning Research, 18:1–52, 2018.

[Mahdisoltani et al. 2013] Farzaneh Mahdisoltani, Joanna Biega, and Fabian M Suchanek. Yago3: a knowledge base from multilingual wikipedias. In CIDR. 2013.

[Ma et al 2018] Embedding models for episodic knowledge graphs. Journal of Web Semantics, 2018.

[Miller 1995] George A Miller. Wordnet: a lexical database for english. Communications of the ACM, 38(11):39–41, 1995.

[Minervini 2017] Minervini, Pasquale, Luca Costabello, Emir Muñoz, Vít Nováček, and Pierre-Yves Vandenbussche. "Regularizing knowledge graph embeddings via equivalence and inversion axioms.”. ECML, 2017.

[Minervini 2020] Minervini, Pasquale, Sebastian Riedel, Pontus Stenetorp, Edward Grefenstette, and Tim Rocktäschel. "Learning Reasoning Strategies in End-to-End Differentiable Proving." ICML, 2020.

[Nguyen et al. 2018] Dai Quoc Nguyen, Tu Dinh Nguyen, Dat Quoc Nguyen, and Dinh Phung. A Novel Embedding Model for Knowledge Base Completion Based on Convolutional Neural Network. NAACL-HLT, 2018.

[Nickel et al. 2016a] Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. A review of relational machine learning for knowledge graphs. Procs of the IEEE, 104(1):11–33, 2016.

[Nickel et al. 2016b] Maximilian Nickel, Lorenzo Rosasco, Tomaso A Poggio, and others. Holographic embeddings of knowledge graphs. AAAI, 1955–1961. 2016.

[Pai and Costabello 2020] Learning Embeddings from Knowledge Graphs With Numeric Edge Attributes. IJCAI, 2021

[Pezeshkpour et al 2019] Pezeshkpour, Pouya, Yifan Tian, and Sameer Singh. Investigating Robustness and Interpretability of Link Prediction via Adversarial Modifications. NAACL, 2019.

[Ren et al. 2020] Ren, Hongyu, Weihua Hu, and Jure Leskovec. "Query2box: Reasoning over knowledge graphs in vector space using box embeddings.”, ICLR 2020

[Rocktäschel et al. 2017] Rocktäschel, Tim, and Sebastian Riedel. "End-to-end differentiable proving." In NIPS 2017.

[Ruffinelli et al 2020] Ruffinelli, Daniel, Samuel Broscheit, and Rainer Gemulla. "You CAN teach an old dog new tricks! on training knowledge graph embeddings." ICLR, 2020.

[Socher et al. 2013] Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. Reasoning with neural tensor networks for knowledge base completion. NIPS, 2013.

[Suchanek et al. 2007] Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. Yago: a core of semantic knowledge. WWW, 2007.

[Sun et al. 2019] Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. Rotate: knowledge graph embedding by relational rotation in complex space. In ICLR, 2019.

[Tabacof & Costabello 2020] Pedro Tabacof and Luca Costabello. Probability Calibration for Knowledge Graph Embedding Models. In ICLR, 2020.

[Toutanova et al. 2015] Kristina Toutanova, Danqi Chen, Patrick Pantel, Hoifung Poon, Pallavi Choudhury, and Michael Gamon. Representing text for joint embedding of text and knowledge bases. EMNLP, 2015.

[Trouillon et al. 2016] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. ICML, 2016.

[Yang et al. 2014] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. arXiv preprint, 2014.

[Zhang et al. 2019] Data Poisoning Attack against Knowledge Graph Embedding, IJCAI 2019.