CS 839: Foundation Models
Knowledge

Fred Sala

University of Wisconsin-Madison

Oct. 26, 2023

Announcements

- **Logistics:**
  - HW2 out (due Nov. 9).
  - Presentation dates: **Nov:** 9,14,16,21,28,30 **Dec:** 5,7
    - Two slots/date (worst-case three might be needed in some special cases)
    - Please sign up! https://docs.google.com/spreadsheets/d/1SqXAtm6VXyofmKh0U3jaH8qg0v6nydnxoptauI8Z_1g/edit#gid=0
- Class roadmap:

Thursday Oct. 26	Knowledge
Tuesday Oct. 31	Scaling & Scaling Laws
Thursday Nov. 2	Security, Privacy, Toxicity
Tuesday Nov. 7	The Future

Outline

- **Knowledge in ML**
  - Short history of KBs/KGs, representing KGs, models for downstream tasks
- **Integrating Knowledge Graphs into FMs**
  - Multimodal-like integration, fusion, RAG, using KGs for hallucination reduction
- **Using FMs for KGs**
  - Crafting KGs from FMs, LLM encoders for KG procedures, etc.

Outline

- **Knowledge in ML**
  - Short history of KBs/KGs, representing KGs, models for downstream tasks
- **Integrating Knowledge Graphs into FMs**
  - Multimodal-like integration, fusion, RAG, using KGs for hallucination reduction
- **Using FMs for KGs**
  - Crafting KGs from FMs, LLM encoders for KG procedures, etc.

Short History of Knowledge Bases/Graphs

Convenient way to represent information about the world

- Classical approach: triplets (head, relation, tail)
  - **Ex:** (Albert Einstein, WinnerOf, Nobel Prize)
- Encode into a **graph**

(Albert Einstein, **BornIn**, German Empire)
(Albert Einstein, **SonOf**, Hermann Einstein)
(Albert Einstein, **GraduateFrom**, University of Zurich)
(Albert Einstein, **WinnerOf**, Nobel Prize in Physics)
(Albert Einstein, **ExpertIn**, Physics)
(Nobel Prize in Physics, **AwardIn**, Physics)
(The theory of relativity, **TheoryOf**, Physics)
(Albert Einstein, **SupervisedBy**, Alfred Kleiner)
(Alfred Kleiner, **ProfessorOf**, University of Zurich)
(The theory of relativity, **ProposedBy**, Albert Einstein)
(Hans Albert Einstein, **SonOf**, Albert Einstein)

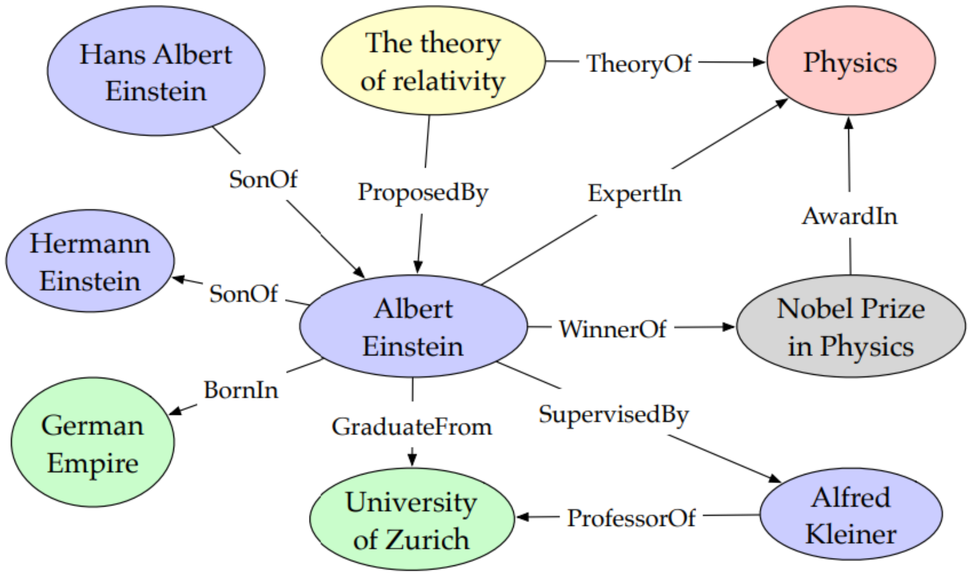

Ji et al `21

Short History of KGs: **Examples**

Convenient way to represent information about the world

- Lots of hand-crafted KBs/KGs out there
- WordNet, YAGO, Freebase, Wikidata

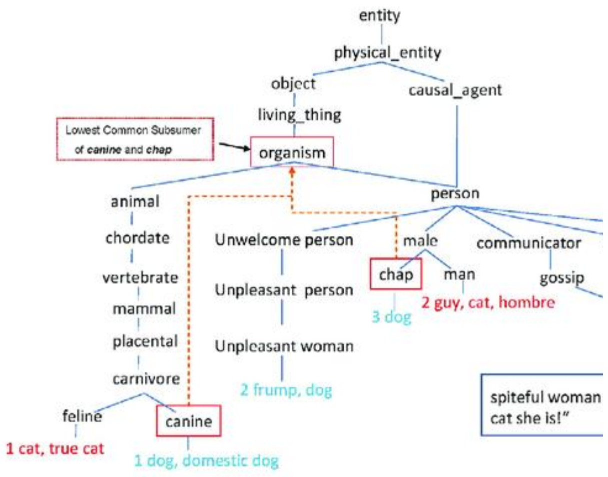

Wei et al ,15

Yago

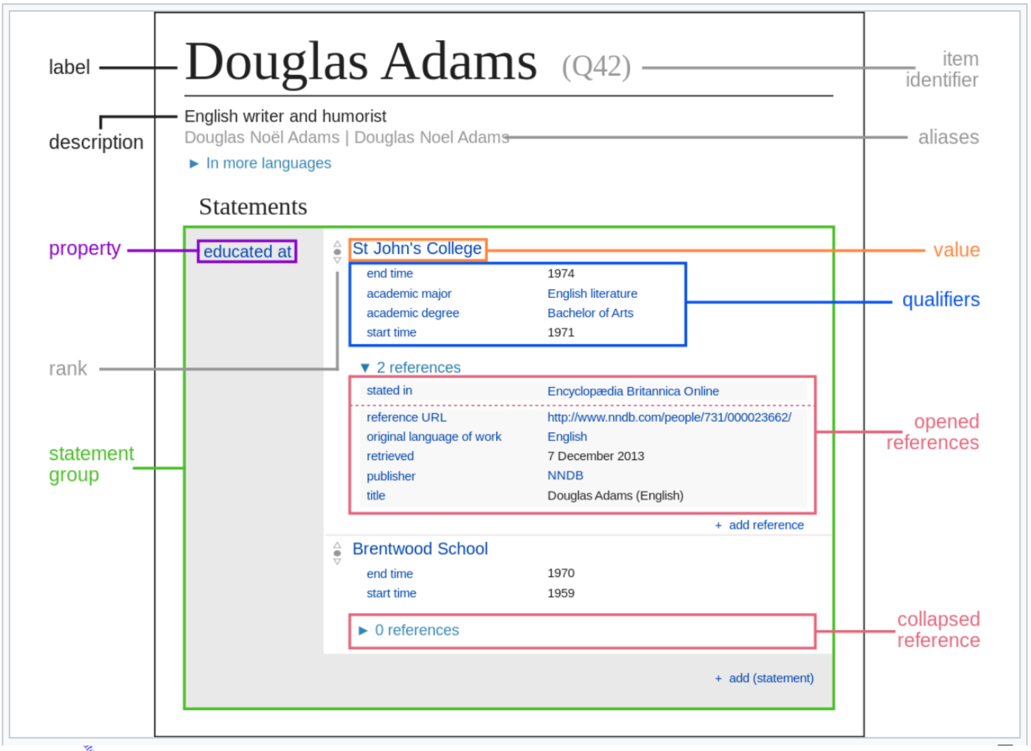

Wikidata

Short History of KGs: **Use**

Reasoning!

- Logical equivalences:

$$(Y, \text{sonOf}, X) \leftarrow (X, \text{hasChild}, Y) \wedge (Y, \text{gender}, \text{Male})$$

- Multi-hop for question-answering or missing facts
  - **Q**: What country was Barack Obama born in?

$$(\text{Obama}, \text{bornIn}, \text{Hawaii}) \wedge (\text{Hawaii}, \text{locatedIn}, \text{US}) \Rightarrow (\text{Obama}, \text{bornin}, \text{US})$$

A: Obama was born in the US.

Short History of KGs: **Embeddings**

KGs are discrete, so we need to represent them in a ML-friendly way

- KG embeddings: a fairly big ML area
- Classic way: embed each entity and relation in R^d in a way that can recover. **TransE** (Bordes et al '13): use loss

$$\mathcal{L} = \sum_{(h,\ell,t)\in S} \sum_{(h',\ell,t')\in S'_{(h,\ell,t)}} \left[\gamma + d(\boldsymbol{h} + \boldsymbol{\ell}, \boldsymbol{t}) - d(\boldsymbol{h'} + \boldsymbol{\ell}, \boldsymbol{t'})\right]_+$$

- Note: can also infer new relations that we didn't have in KG (good, since all KGs are **incomplete**!)

Short History of KGs: **Embeddings**

KG embeddings: a fairly big ML area

- Many **variations**: instead of relationship a vector to be added, can parametrize some transformation instead

$$f_r(h, t) = h^T \mathrm{diag}(M_r) t$$

Yang et al `14

- Can use some **other spaces** to avoid Euclidean distance-based scoring
  - Complex vector spaces C
  - Quaternions
  - Non-Euclidea spaces

Quaternion multiplication table

↓×→	1	i	j	k
1	1	i	j	k
i	i	−1	k	−j
j	j	−k	−1	i
k	k	j	−i	−1

Short History of KGs: **Non-Euclid. Embeddings**

Can use some **other spaces** to avoid Euclidean distance-based scoring: hyperbolic space

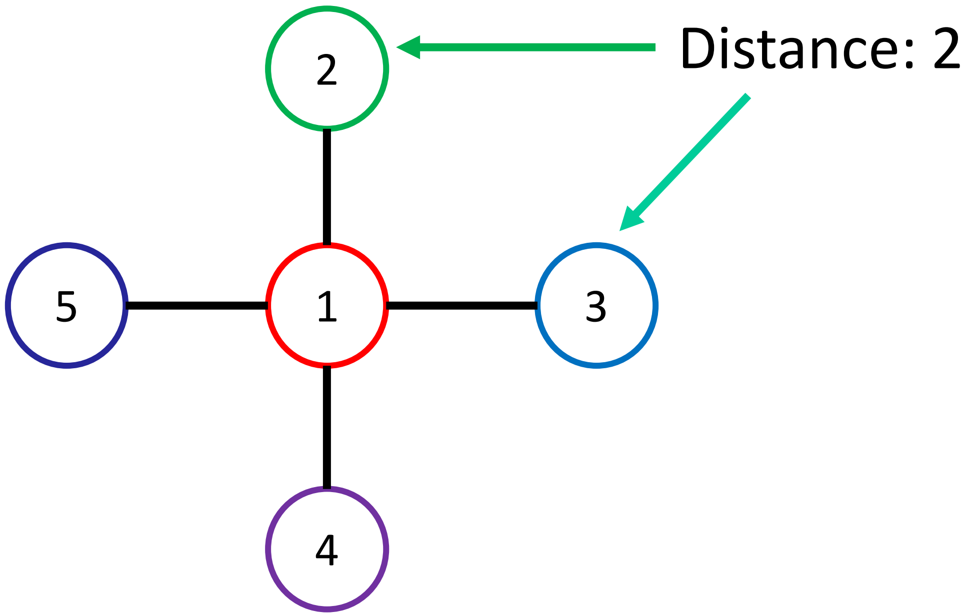

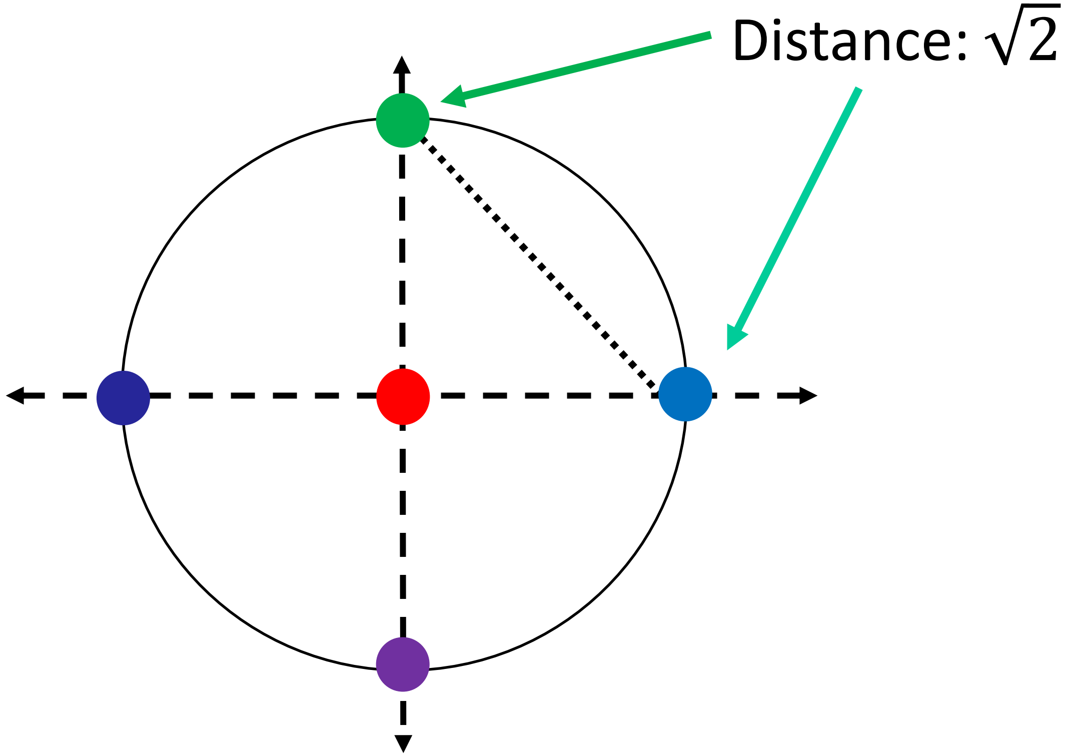

Distortion!

$$\frac{\max_{x,y\in V} \frac{d_S(f(x),f(y))}{d_G(x,y)}}{\min_{x,y\in V} \frac{d_S(f(x),f(y))}{d_G(x,y)}} = \frac{\text{Biggest Stretch}}{\text{Worst Compression}}$$

Short History of KGs: **Hyperbolic Embeddings**

- Poincaré **model** of hyperbolic space

- Connection to **hierarchical distance**:
  - Hyperbolic distance:

$$d_H(x, y) = \operatorname{acosh}\left(1 + 2\frac{\|x - y\|^2}{(1 - \|x\|^2)(1 - \|y\|^2)}\right)$$

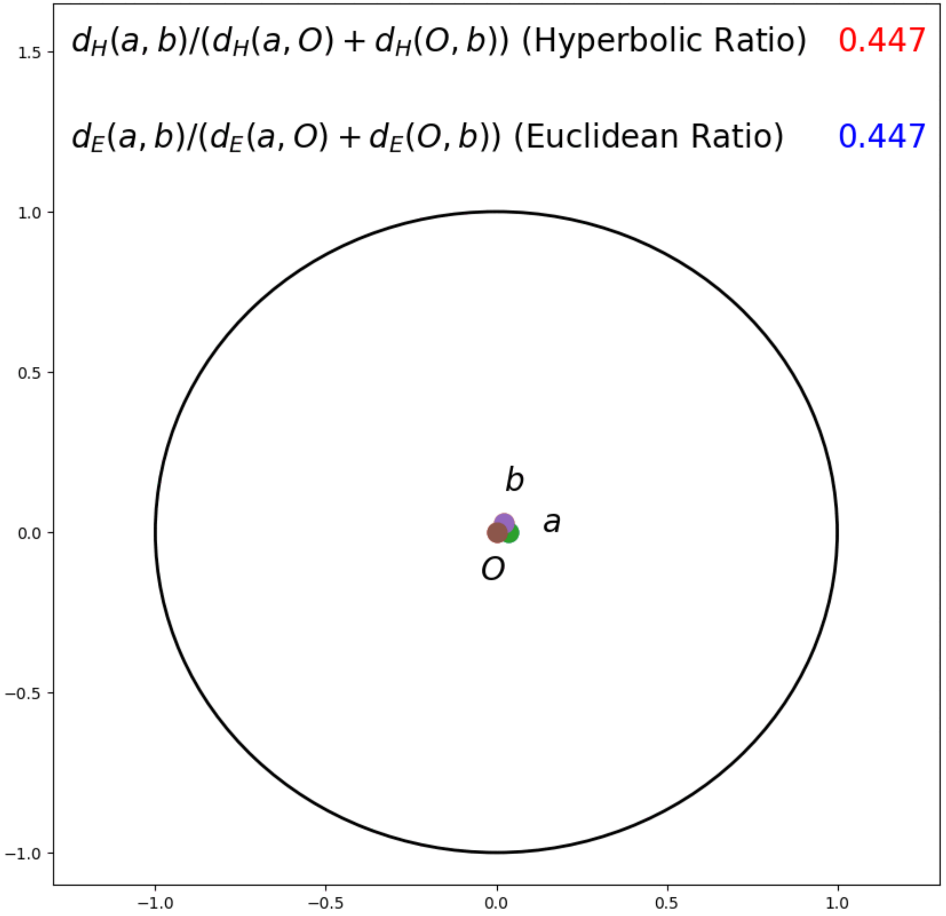

Short History of KGs: **Hyperbolic Embeddings**

- Optimize with non-Euclidean variants of SGD

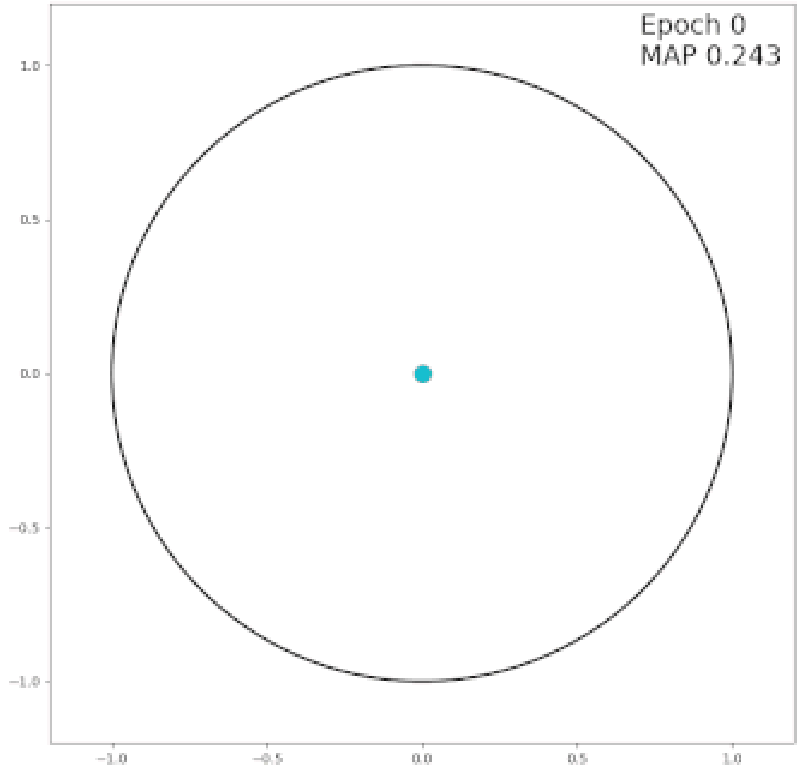

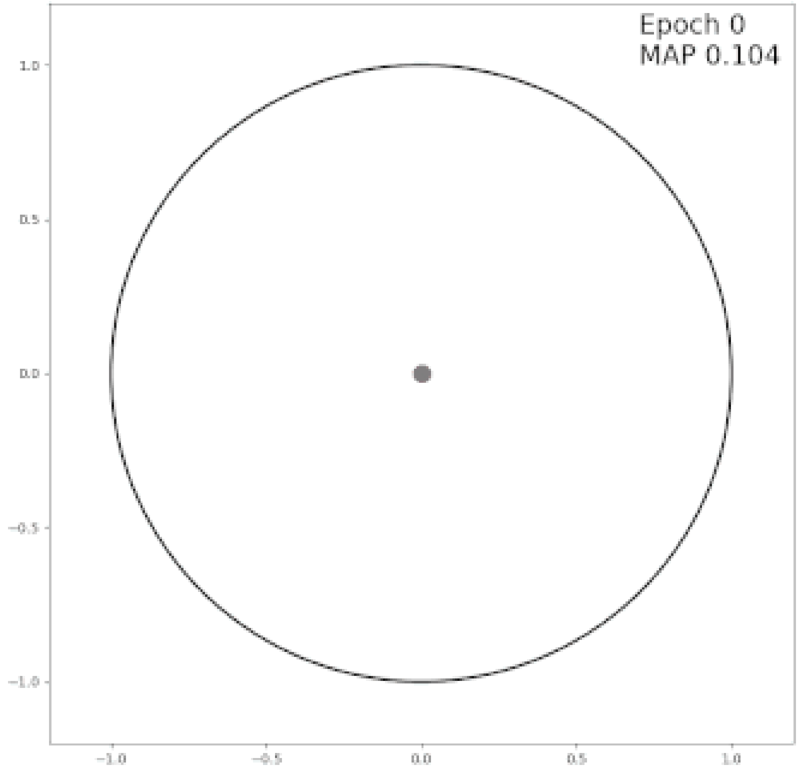

Short History of KGs: **GNNs**

Post-embeddings, use for downstream task

- Can combine representation learning + task
- GCNs:

$$f(X, A) = \mathrm{softmax}(A\sigma(AXW^{(0)})W^{(1)})$$

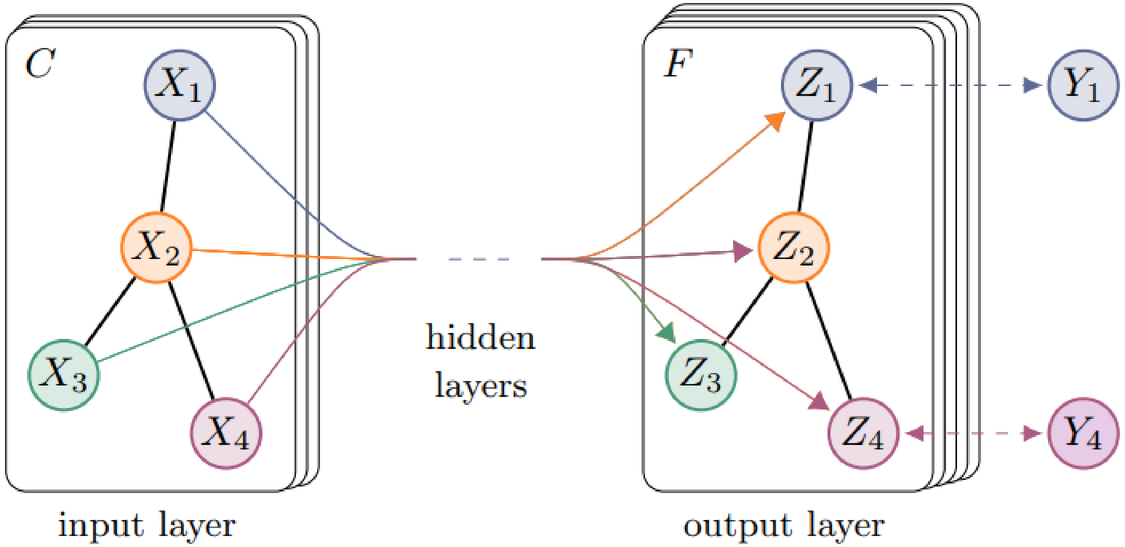

Adjacency Matrix

Layer 1 Weights

Layer 2 Weights

Kipf and Welling '16

Short History of KGs: **Transformers**

Can also use transformers as well

- Ex: **KG-BERT**

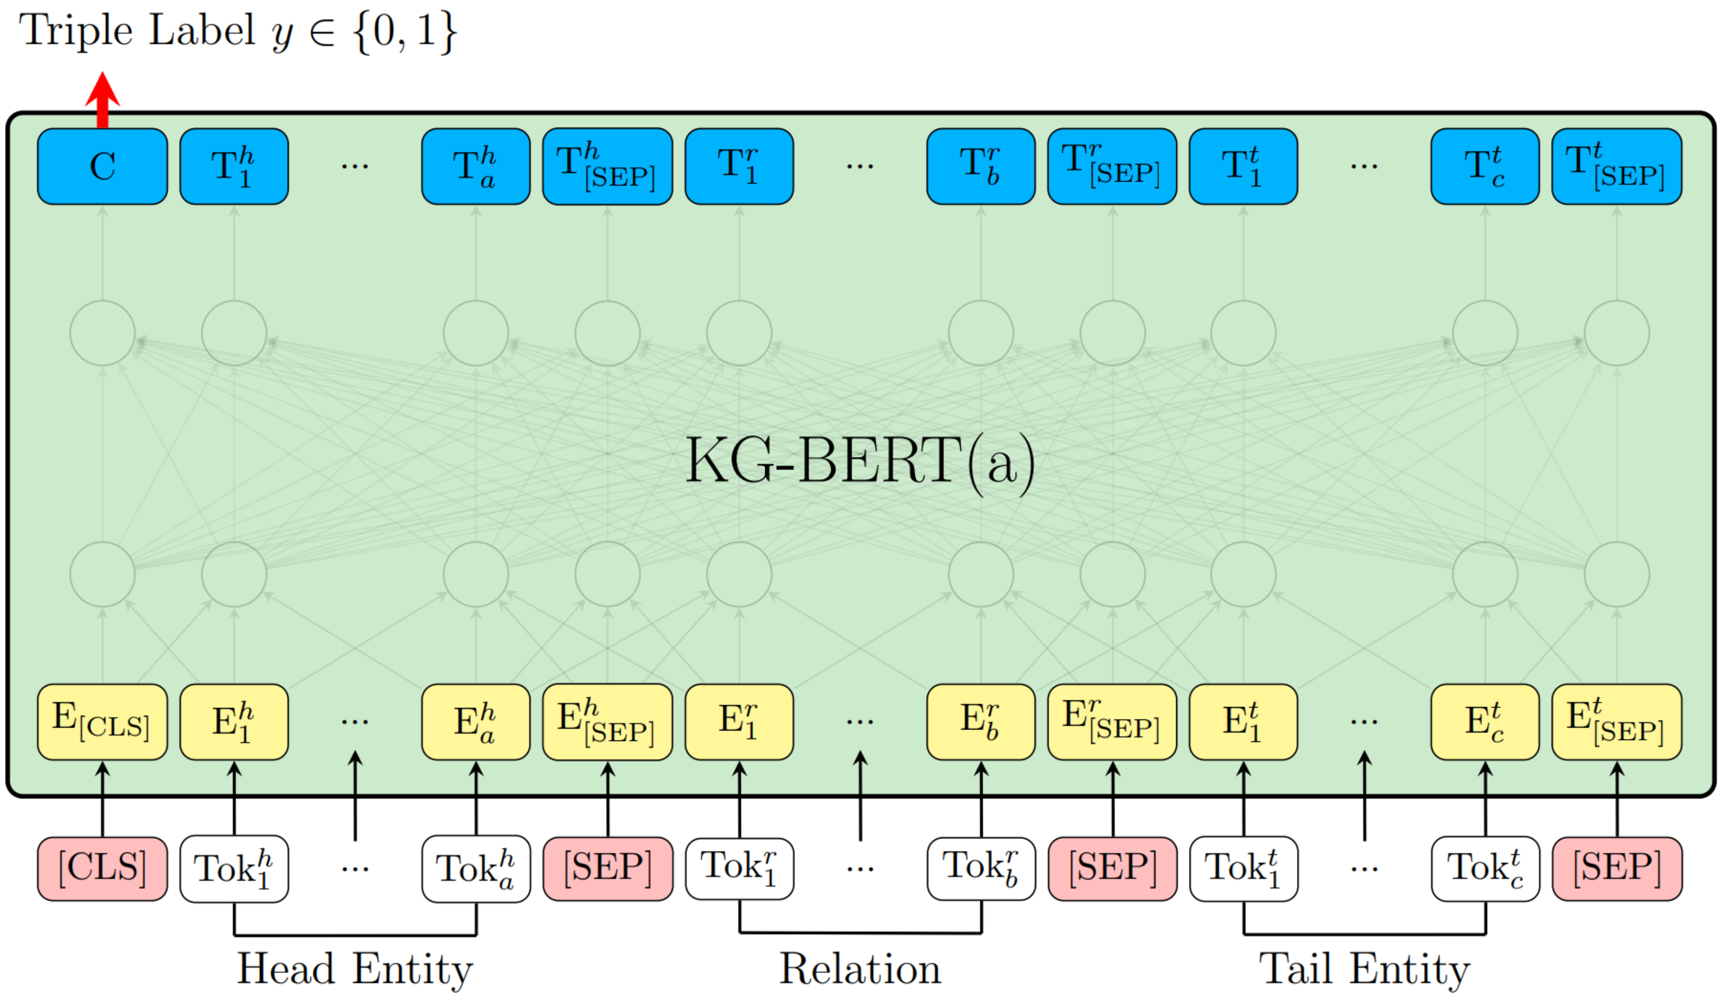

Yao et al '19

Break & Questions

Outline

- **Knowledge in ML**
  - Short history of KBs/KGs, representing KGs, models for downstream tasks
- **Integrating Knowledge Graphs into FMs**
  - Multimodal-like integration, fusion, RAG, using KGs for hallucination reduction
- **Using FMs for KGs**
  - Crafting KGs from FMs, LLM encoders for KG procedures, etc.

Mixing LLMs and KGs: Why?

Both approaches have downsides by themselves

KGs:

- Need lots of supervised data (often manual) to create
- Fixed schemas for questions
- Incomplete and noisy

FMs:

- Hallucinate
- Not aware of stale facts + hard to update

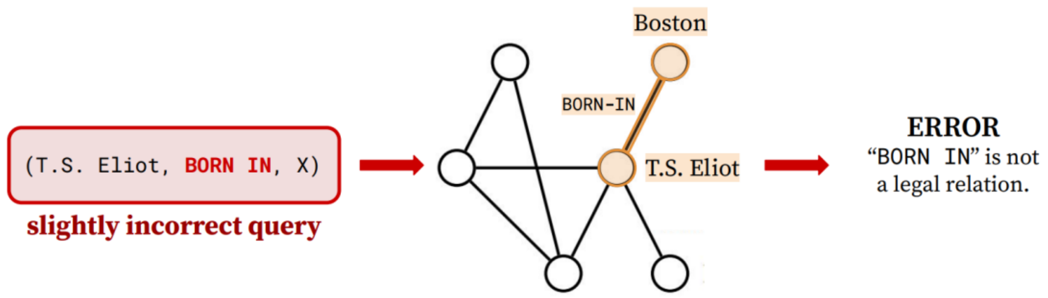

Princeton CS579

Mixing LLMs and KGs: **Approaches**

A nice categorization:

- Pan et al '23, "Unifying Large Language Models and Knowledge Graphs: A Roadmap"

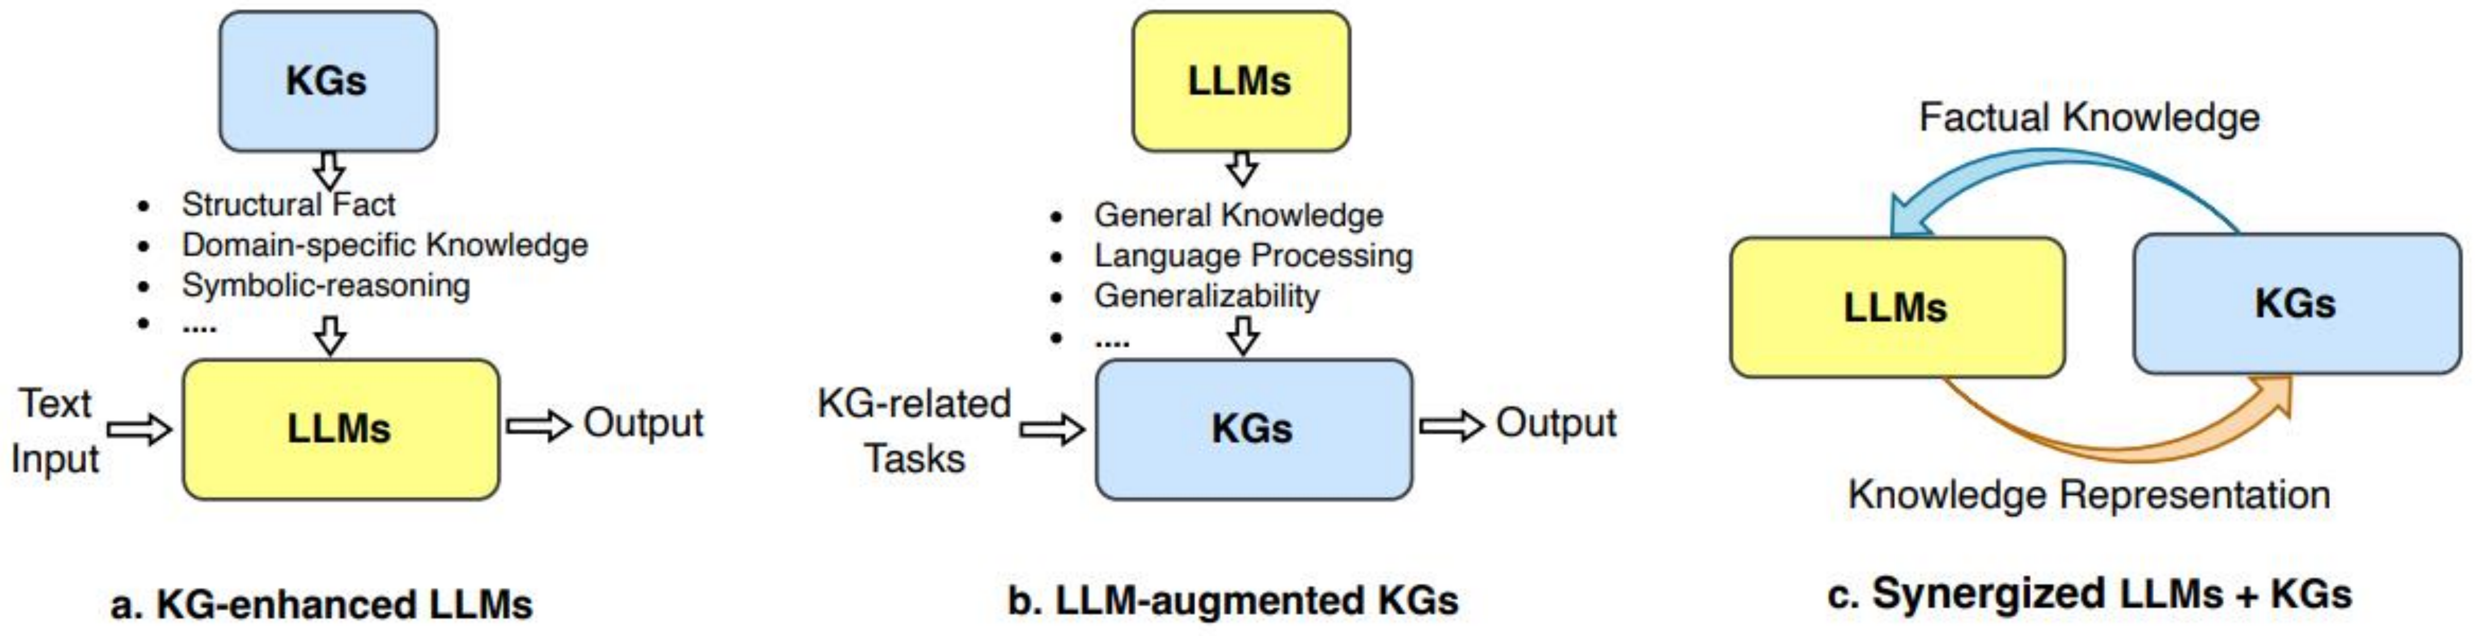

KG-Enhanced LLMs: Pretraining

Goal: create more structured information in LLMs

- In pretraining: add information from KG

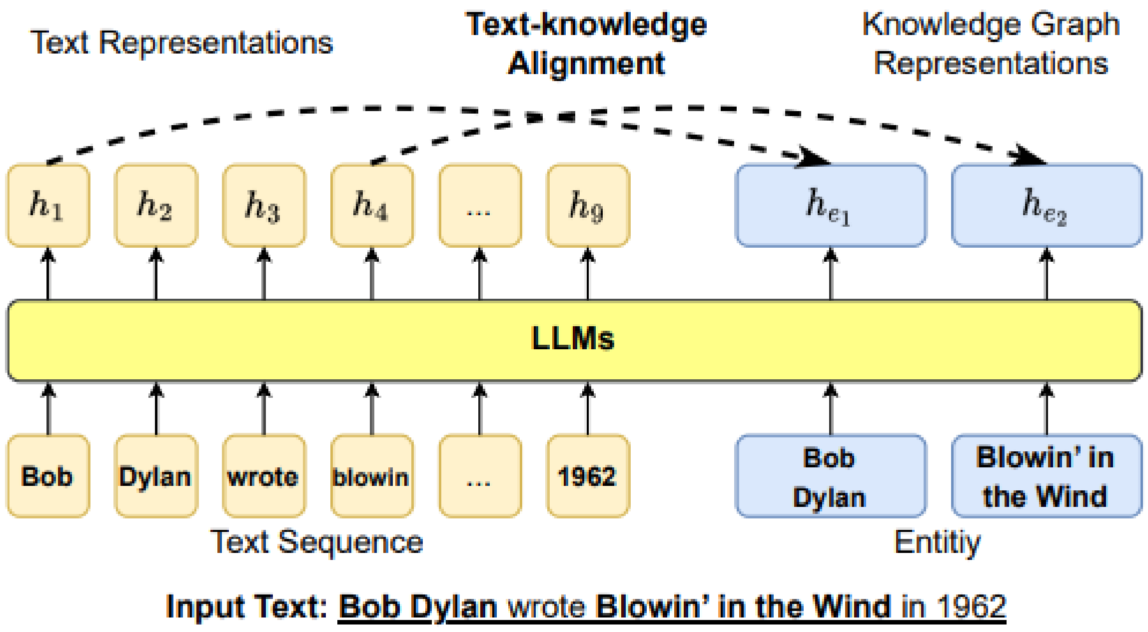

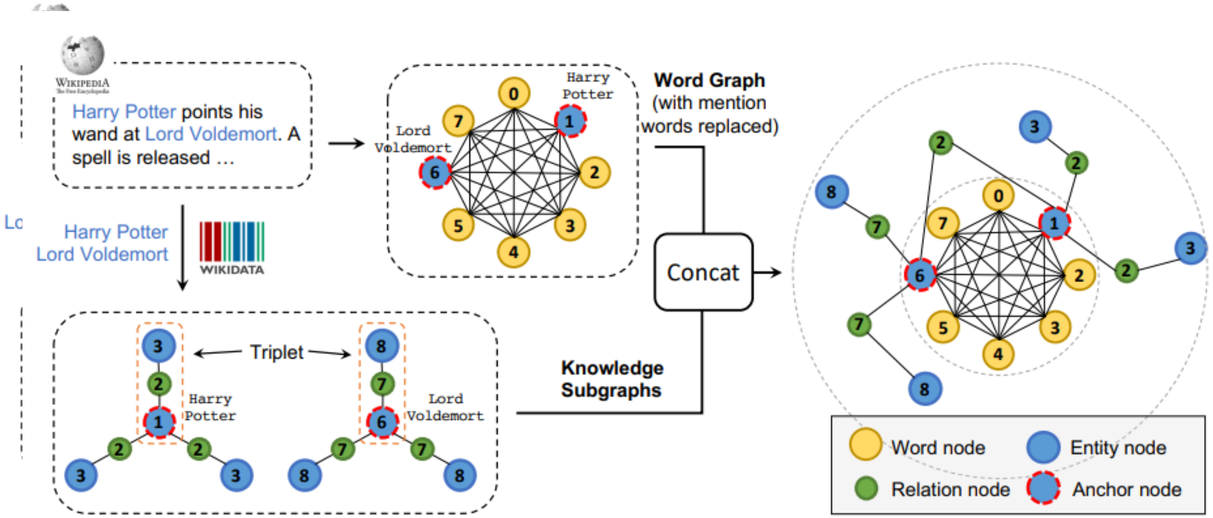

Sun et al '20

KG-Enhanced LLMs: **Fusion**

Goal: create more structured information in LLMs

- Note: similar to **multimodal models**, can do **fusion** earlier or later

- **Ex**: allow text processing,
  combine with a later module

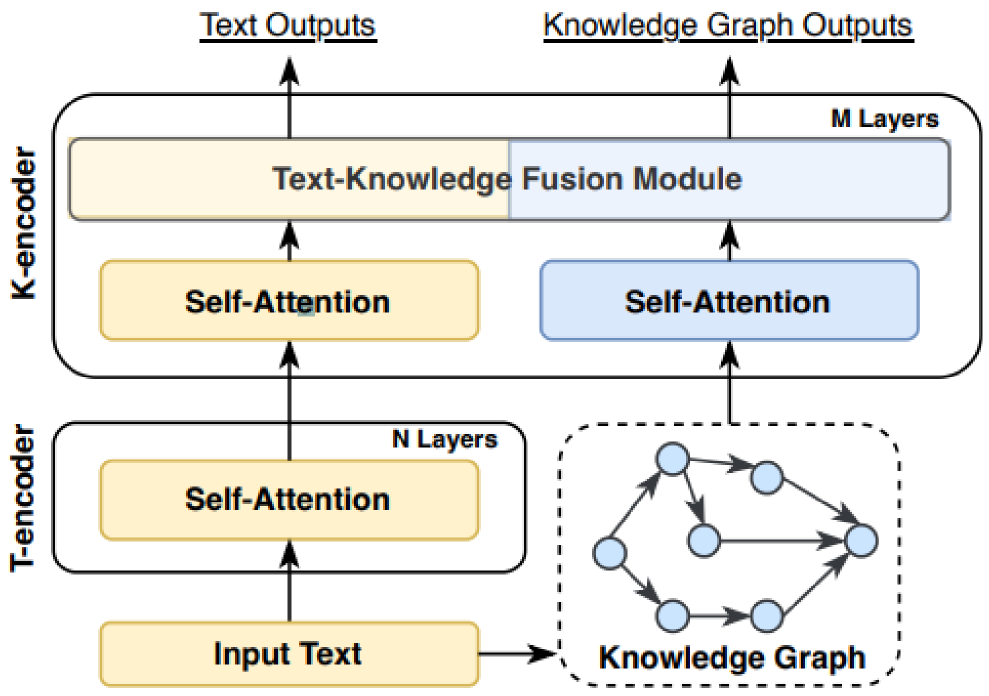

KG-Enhanced LLMs: **RAG**

Retrieval-Augmented Generation (RAG)

- Not strictly just for knowledge graphs
- Search for relevant documents, use results as part of context for generator

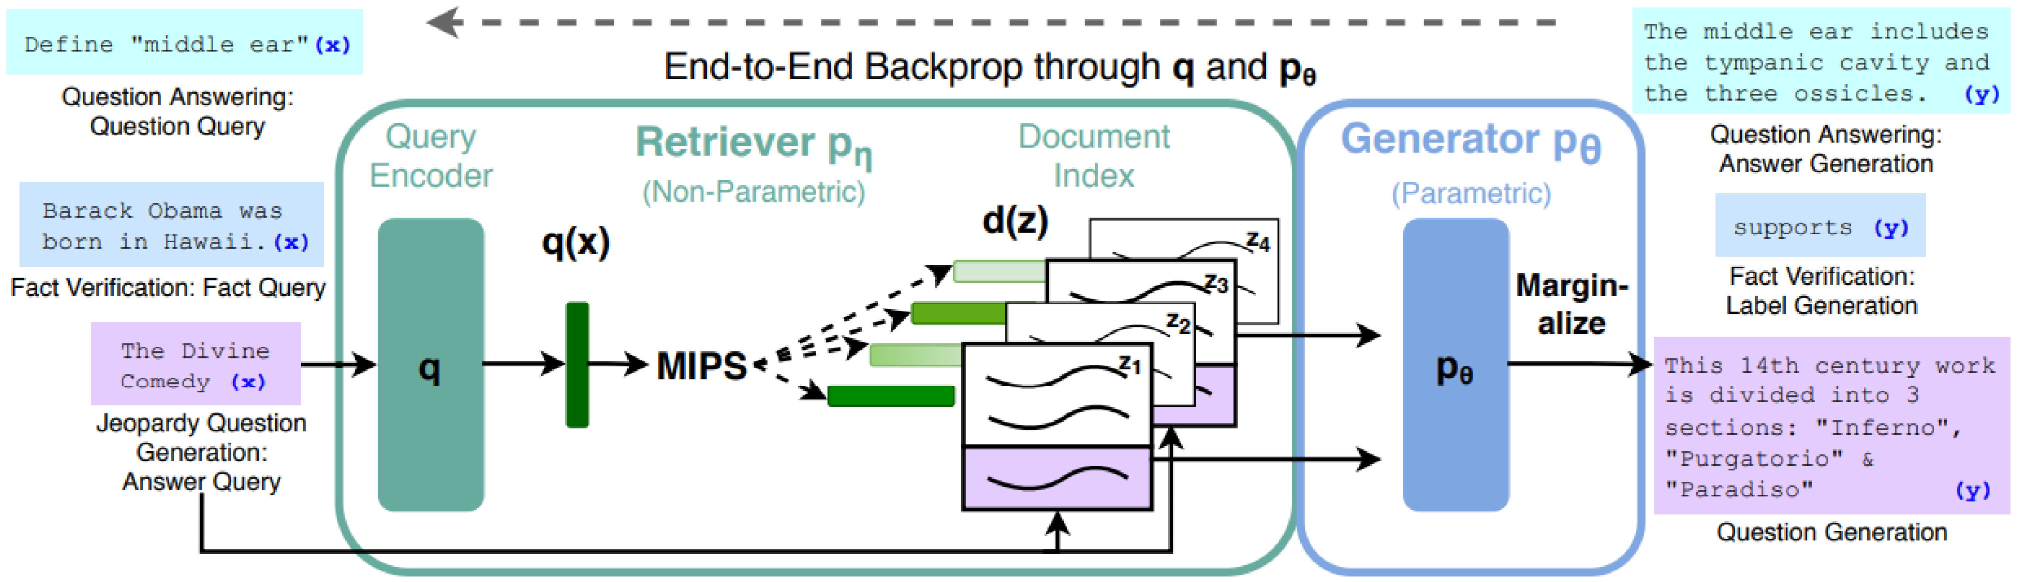

Lewis et al `21

RAG Components: **Retriever Side**

Let's walk through each component in RAG

- Query encoder q(x)
- Document encoder d(z)

$$p_\eta(z|x) \propto \exp\left(\mathbf{d}(z)^\top \mathbf{q}(x)\right) \qquad \mathbf{d}(z) = \mathrm{BERT}_d(z), \quad \mathbf{q}(x) = \mathrm{BERT}_q(x)$$

Retriever **Documents** **Query**

- Obtain k documents z with highest prob. using **MIPS**

RAG Components: **Generator Side**

Let’s walk through each component in RAG

- Generator: sample y conditioned on x, zs
- Any sequence-to-sequence model

$$p_\theta(y_i | x, z, y_{1:i-1})$$

- Obtain k documents z with highest prob. using **MIPS**
- **Note:** backprop through whole thing, but easier to keep the document encoder fixed.

RAG **Results**

How does it do compared to non-augmented LLMs?

Task	Input	Model	Generation
MS-MARCO	define middle ear	BART	[?]The middle ear is the part of the ear between the middle ear and the nose.
		RAG-T	The middle ear is the portion of the ear internal to the eardrum.
		RAG-S	The middle ear includes the tympanic cavity and the three ossicles.
	what currency needed in scotland	BART	The currency needed in Scotland is Pound sterling.
		RAG-T	Pound is the currency needed in Scotland.
		RAG-S	The currency needed in Scotland is the pound sterling.
Jeopardy Question Gener -ation	Washington	BART	[?]This state has the largest number of counties in the U.S.
		RAG-T	It's the only U.S. state named for a U.S. president
		RAG-S	It's the state where you'll find Mount Rainier National Park
	The Divine Comedy	BART	[*]This epic poem by Dante is divided into 3 parts: the Inferno, the Purgatorio & the Purgatorio
		RAG-T	Dante's "Inferno" is the first part of this epic poem
		RAG-S	This 14th century work is divided into 3 sections: "Inferno", "Purgatorio" & "Paradiso"

KGs to **Reduce Hallucination**

Recall that LLMs have a tendency to **hallucinate**

- **Example:** summarizing non-existent article

PA

summarise this article https://www.nytimes.com/2023/03/11/technology/chatgpt-prompts-to-avoid-content-filters.html

The article discusses how ChatGPT, an AI language model, is being used to circumvent content filters that are designed to detect and block certain types of content. By generating prompts that are less likely to trigger these filters, ChatGPT can be used to create content that might otherwise be censored. While this technology could be useful for those seeking to express themselves freely, it also raises concerns about the potential misuse of AI-generated content, as well as the difficulty of regulating such content in the future. The article concludes by noting that there is still much to learn about the impact of AI language models like ChatGPT on the way we communicate online.

Wiki

KGs to **Reduce Hallucination**

Reduce it by leveraging existing KG

- Example: Neural Path Hunter (Dziri et al '21)

- Produce a path on KG from LLM outputs
- Compute likelihood of path.
- Throw out paths that are unlikely.

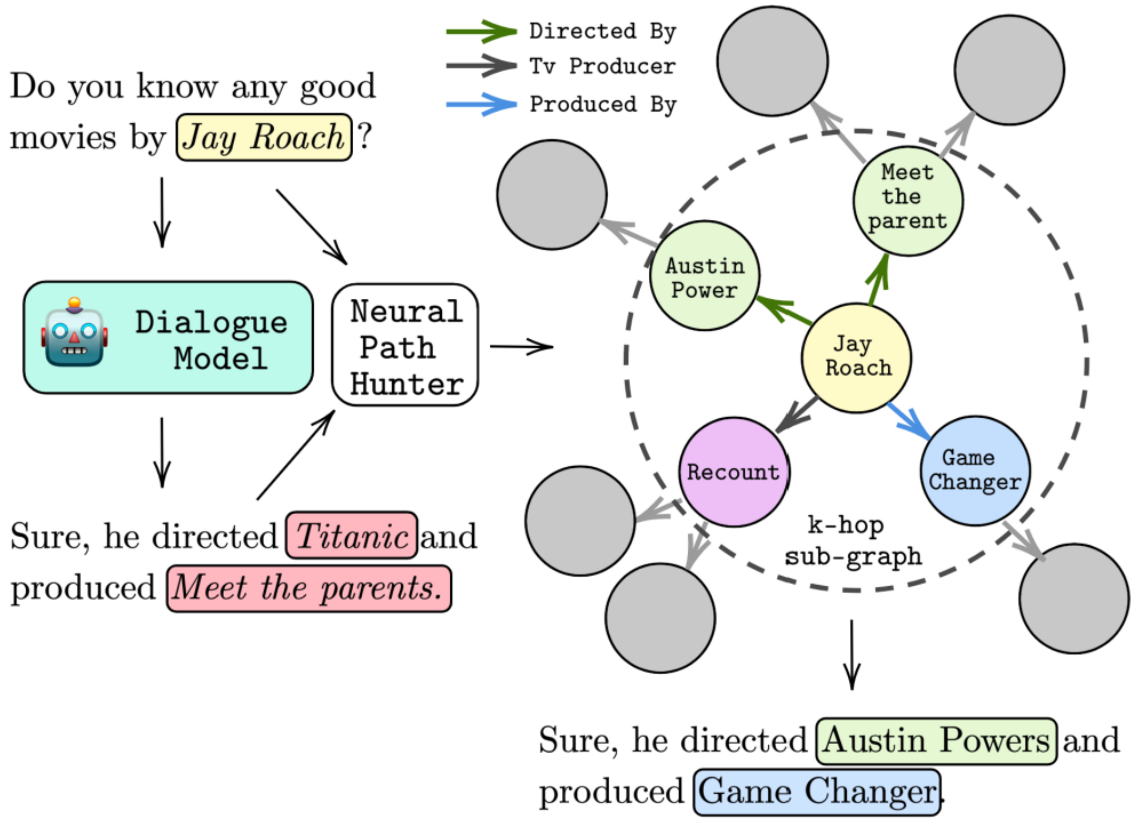

Using KGs to **Evaluate LLMs**

Recall our efforts to evaluate LLMs

- LAMA: LAnguage Model Analysis
- Evaluate **factual** and **commonsense** knowledge

- Transform KG triplets into queries to evaluate the LM

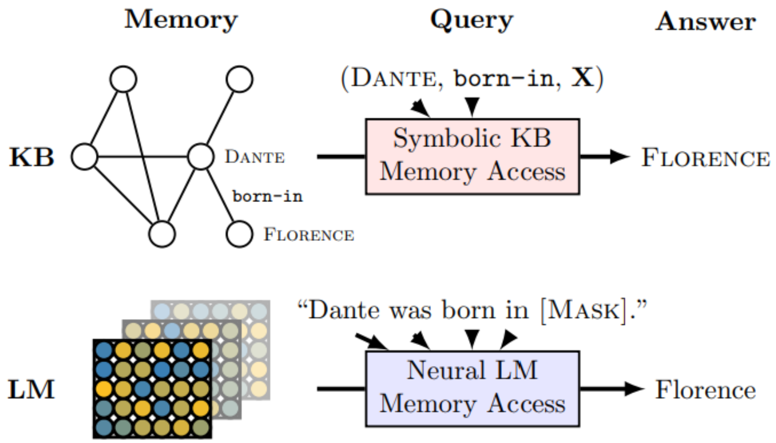

Petroni et al '19

Break & Questions

Outline

- **Knowledge in ML**
  - Short history of KBs/KGs, representing KGs, models for downstream tasks
- **Integrating Knowledge Graphs into FMs**
  - Multimodal-like integration, fusion, RAG, using KGs for hallucination reduction
- **Using FMs for KGs**
  - Crafting KGs from FMs, LLM encoders for KG procedures, etc.

LLMs to help KGs

Other direction: use LLMs to improve KGs.

Recall the downsides of KGs: somewhat inflexible and generally incomplete.

- Completing them: use a mixed model like KGBERT
  - Predict that a triplet is valid or not
- Generating entities for KGs:
  - Prompt an LLM with (h,r,?) and use output to define tail entity.

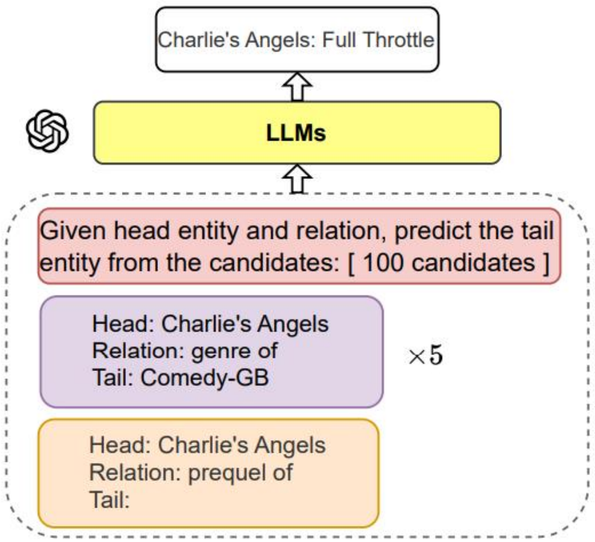

Building KGs From LLMs

Directly **extract** triplets to build KB/KG

- Use attention weights
- Decode directly into the graph structure
- Caveat: might be noisy!

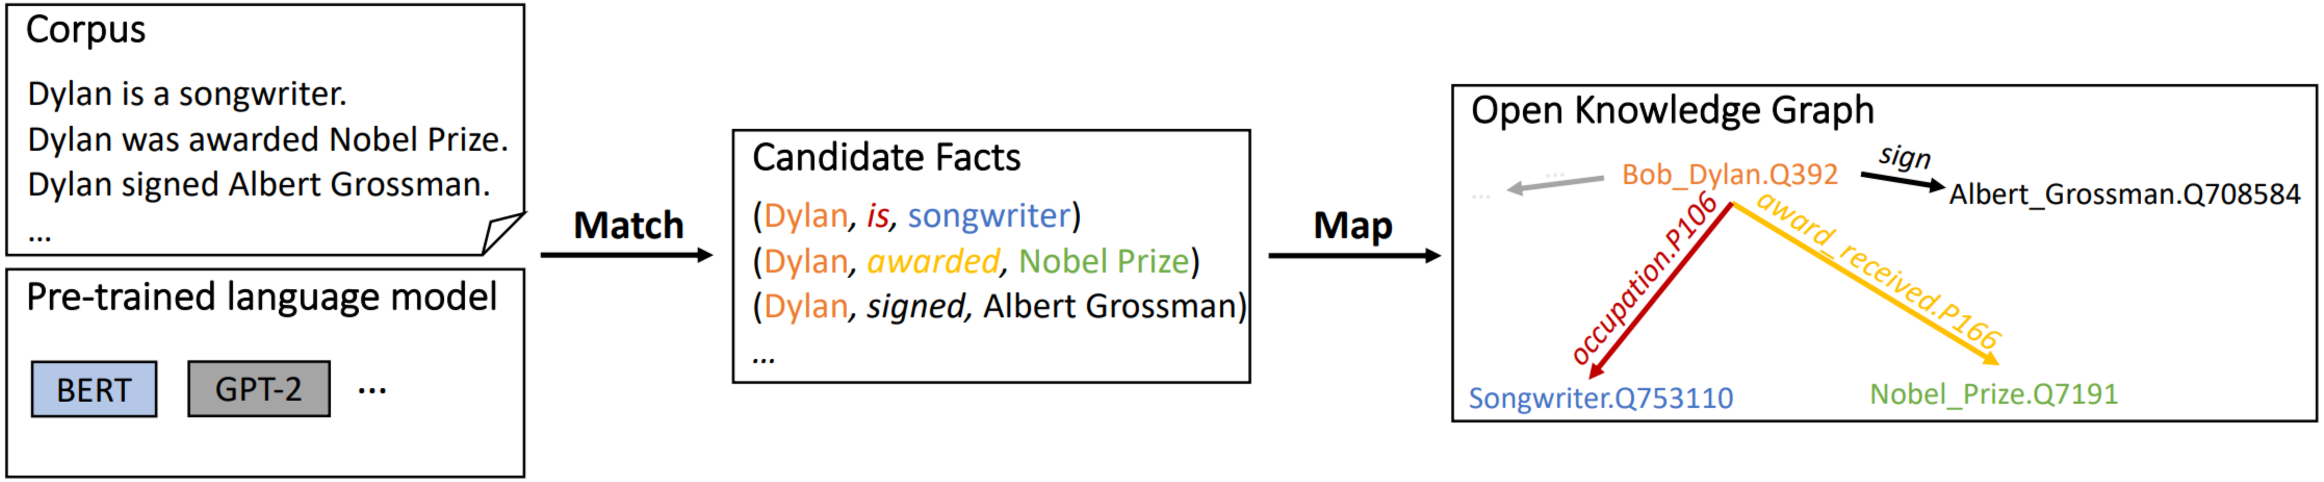

Wang et al '21

Bibliography

- Ji et al '21: Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, Philip S. Yu, "A Survey on Knowledge Graphs: Representation, Acquisition and Applications" (https://arxiv.org/abs/2002.00388)
- Wei et al '15: Tingting Wei, Yonghe Lu, Xuanjie Li, Jinglun Liu, "Web page segmentation based on the hough transform and vision cues" (https://ieeexplore.ieee.org/document/7415394)
- Bordes et al '13: Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, Oksana Yakhnenko, "Translating Embeddings for Modeling Multi-relational Data" (https://papers.nips.cc/paper_files/paper/2013/hash/1cecc7a77928ca8133fa24680a88d2f9-Abstract.html)
- Yang et al '14: Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, Li Deng, "Embedding Entities and Relations for Learning and Inference in Knowledge Bases" (https://arxiv.org/abs/1412.6575)
- Kipf and Welling '16: Thomas N. Kipf, Max Welling, "Semi-Supervised Classification with Graph Convolutional Networks" (https://arxiv.org/abs/1609.02907)
- Yao et al '19: Liang Yao, Chengsheng Mao, Yuan Luo, "KG-BERT: BERT for Knowledge Graph Completion" (https://arxiv.org/abs/1909.03193)
- Princeton CS579G: https://www.cs.princeton.edu/courses/archive/fall22/cos597G/lectures/lec10.pdf
- Pan et al '23: Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiapu Wang, Xindong Wu, "Unifying Large Language Models and Knowledge Graphs: A Roadmap" (https://arxiv.org/abs/2306.08302)
- Sun et al '20: Tianxiang Sun, Yunfan Shao, Xipeng Qiu, Qipeng Guo, Yaru Hu, Xuanjing Huang, Zheng Zhang, "CoLAKE: Contextualized Language and Knowledge Embedding" (https://aclanthology.org/2020.coling-main.327.pdf)
- Lewis et al '21: Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, Douwe Kiela, "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks" (https://arxiv.org/abs/2005.11401)
- Dziri et al '21: Nouha Dziri, Andrea Madotto, Osmar Zaiane, Avishek Joey Bose, Neural Path Hunter: Reducing Hallucination" (https://arxiv.org/pdf/2104.08455.pdf)
- Petroni et al '19: Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H. Miller, Sebastian Riedel, "Language Models as Knowledge Bases?" (https://aclanthology.org/D19-1250.pdf)
- Wang et al '21: Chenguang Wang, Xiao Liu, Dawn Song, "Language Models are Open Knowledge Graphs“ (https://openreview.net/forum?id=aRTRjVPkm-)

Thank You!