



Knowledge Graph Foundation Models for Neuro-Symbolic Reasoning

Zequn Sun

Nanjing University, China

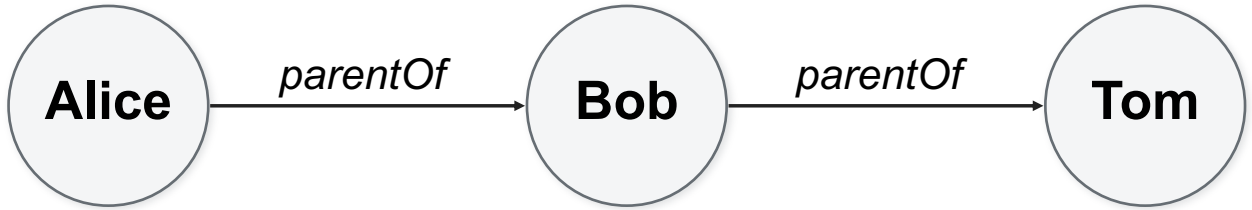
Agenda

1. **Introduction**
2. Knowledge Graph Reasoning
3. Knowledge Graph Entity Alignment
4. Knowledge Graph Retrieval
5. Conclusion and Future Work

Neuro-Symbolic Reasoning and KGs

From explicit logical rules to reusable graph reasoning

Shared facts



Logical deduction

RULE

parentOf(x, y) ∧ *parentOf*(y, z) ⇒ *grandparentOf*(x, z)

DERIVED FACT

grandparentOf(Alice, Tom)

Follows from the facts and rule.

Learned graph reasoning

QUERY

(Alice, *grandparentOf*, ?)

RANKED ANSWER

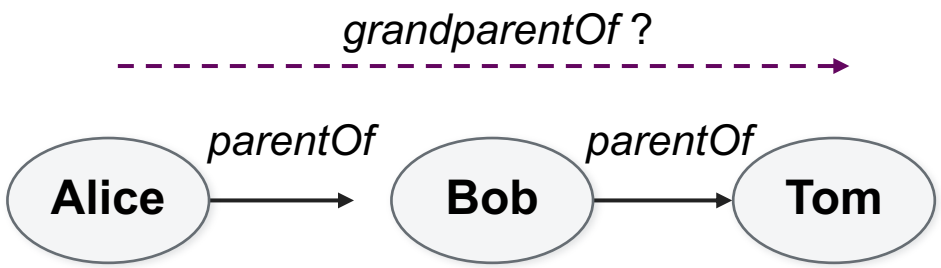
Tom, …

A prediction, without a proof guarantee.

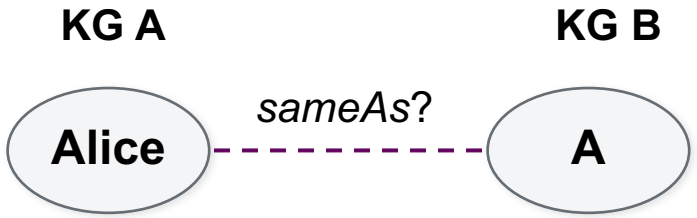
Incomplete and Distributed Knowledge

Real KGs are incomplete, distributed, and heterogeneous

Within one knowledge graph



Across knowledge graphs



Missing facts

A valid grandparent link may be missing, even when its evidence is present.

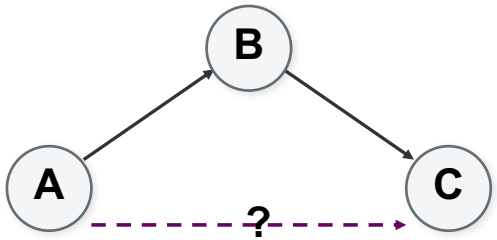
Different identifiers

The same entity can have different IDs and neighborhoods in another KG.

Three Tasks, One Transfer Challenge

Three tasks use different inputs and produce different outputs

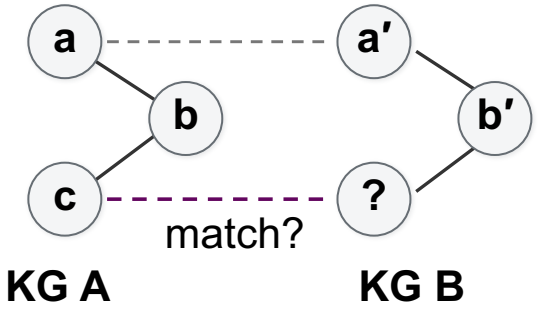
Reasoning



Predict missing facts

KG + query → ranked facts

Entity alignment

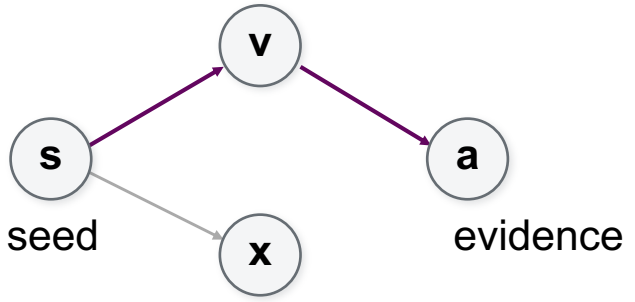


Match entities

Two KGs → ranked matches

Retrieval

Question: who is relevant?



Retrieve evidence

Question + KG → evidence

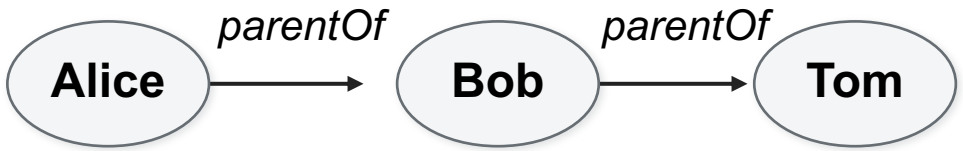
Agenda

1. Introduction
2. **Knowledge Graph Reasoning**
3. Knowledge Graph Entity Alignment
4. Knowledge Graph Retrieval
5. Conclusion and Future Work

Link Prediction

A concrete query and the ranked output

Observed facts



Query

(Alice, *grandparentOf*, ?)

Candidate ranking

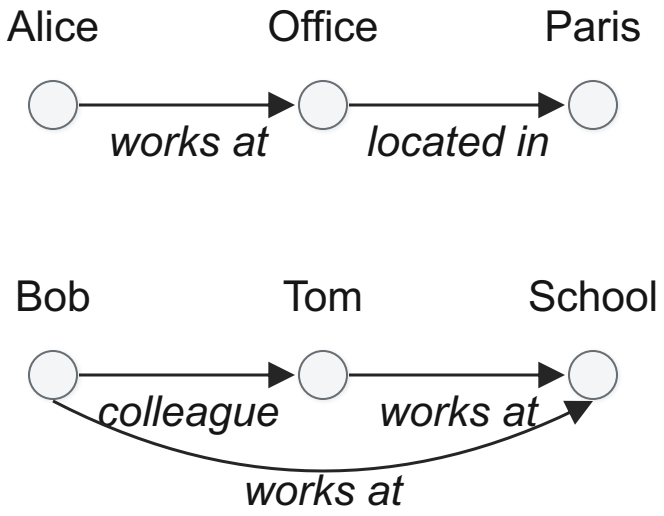
Rank	Candidate
1	**Tom**
2	Lee
3	…

Ranking quality: MRR / Hits@K

Transfer Settings

Pretrain on source graphs; reuse the model on a target graph

Pretraining graphs

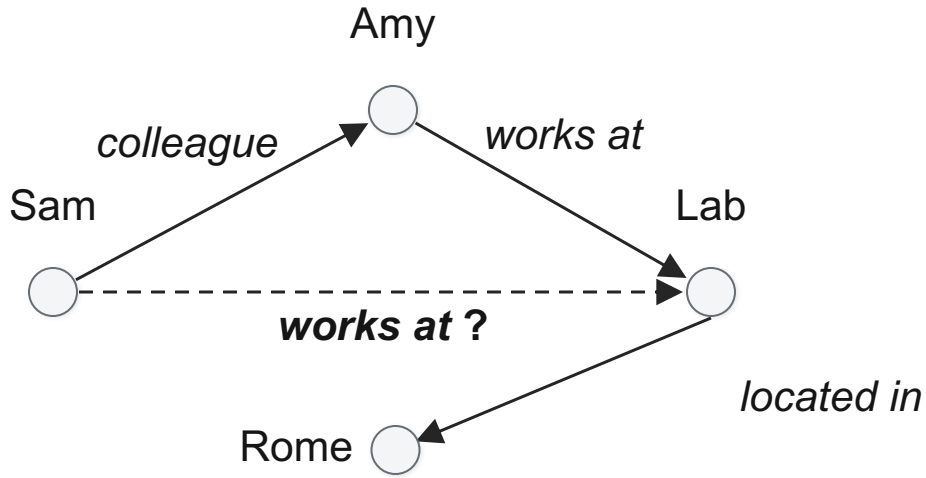


Shared model

Trained on source graphs

reuse

A new target graph



Target: new entities and/or relations

Dashed: candidate link

Query: (Sam, *works at*, ?)

Zero-shot: no target fine-tuning.

Inductive Graph Reasoning

Inductive reasoning as the background for graph foundation models

NBFNet: entity paths



InGram: relation interactions



Query-conditioned path aggregation

Start at the query head; aggregate paths to score candidates.

Shared operations handle unseen entities.

Relation graph as shared context

Relation interactions guide entity-level reasoning.

New relations use target-graph structure.

[1] Zhaocheng Zhu, Zuobai Zhang, Louis-Pascal Xhonneux, Jian Tang. Neural Bellman-Ford Networks: A General Graph Neural Network Framework for Link Prediction. NeurIPS 2021.
[2] Jaejun Lee, Chanyoung Chung, Joyce Jiyoung Whang. InGram: Inductive Knowledge Graph Embedding via Relation Graphs. ICML 2023.

Contexts for Transfer

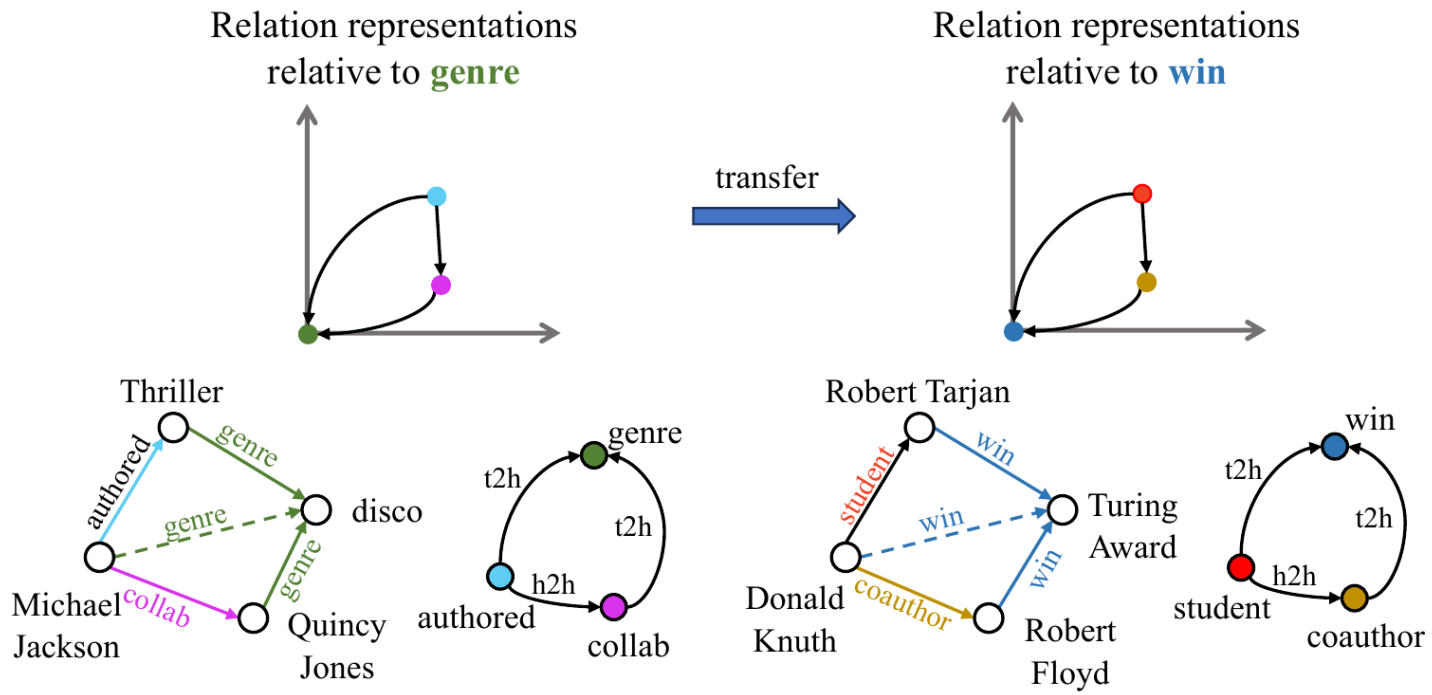
Reasoning models differ in transferable context and target inputs

Model	Transfer signal	Target graph	Extra context
ULTRA [1]	Relation interactions	Graph structure	None
MOTIF [2]	Higher-order motifs	Graph structure	None
ProLINK [3]	Generated type priors	Graph structure	LLM prompts
SEMMA [4]	Textual relation graph	Graph structure	Relation text
KG-ICL [5]	Example graph prompts	Graph structure	Graph examples

[1] Mikhail Galkin, Xinyu Yuan, Hesham Mostafa, Jian Tang, Zhaocheng Zhu. Towards Foundation Models for Knowledge Graph Reasoning. ICLR 2024.
[2] Xingyue Huang, Pablo Barceló, Michael M. Bronstein, İsmail İlkan Ceylan, Mikhail Galkin, Juan L. Reutter, Miguel Romero Orth. How Expressive are Knowledge Graph Foundation Models? ICML 2025.
[3] Kai Wang, Yuwei Xu, Zhiyong Wu, Siqiang Luo. LLM as Prompter: Low-resource Inductive Reasoning on Arbitrary Knowledge Graphs. Findings of ACL 2024.
[4] Arvindh Arun, Sumit Kumar, Mojtaba Nayyeri, Bo Xiong, Ponnurangam Kumaraguru, Antonio Vergari, Steffen Staab. SEMMA: A Semantic Aware Knowledge Graph Foundation Model. EMNLP 2025.
[5] Yuanning Cui, Zequn Sun, Wei Hu. A Prompt-Based Knowledge Graph Foundation Model for Universal In-Context Reasoning. NeurIPS 2024.

ULTRA: Relation Interaction Graph

ULTRA lifts entity-level interactions to relations



Relation interactions

Four interaction types from shared heads or tails.

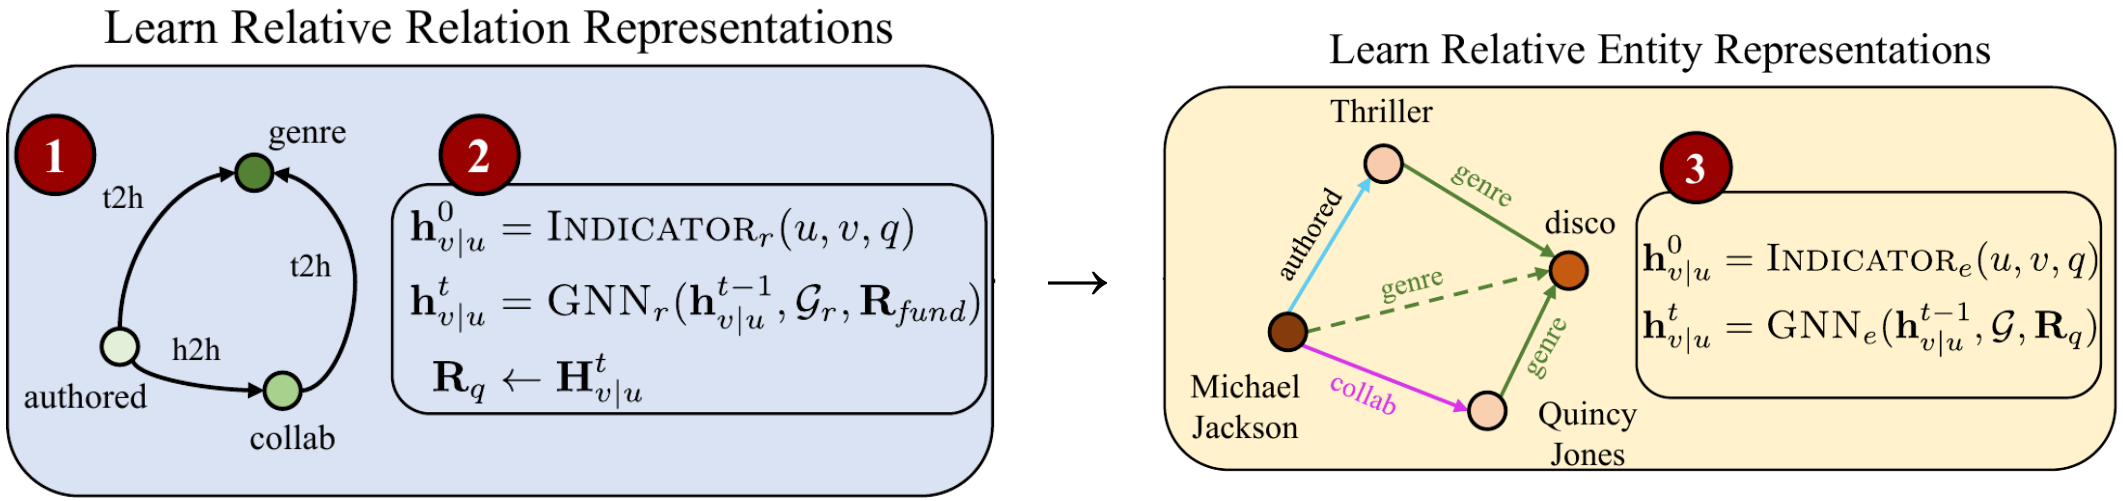
Transfer

Reuse relation patterns without shared entity IDs.

ULTRA: Query-Conditioned Reasoning

Two stages connect a relation query to candidate entities

Query: (head entity, *genre*, ?)



1 Relation encoder

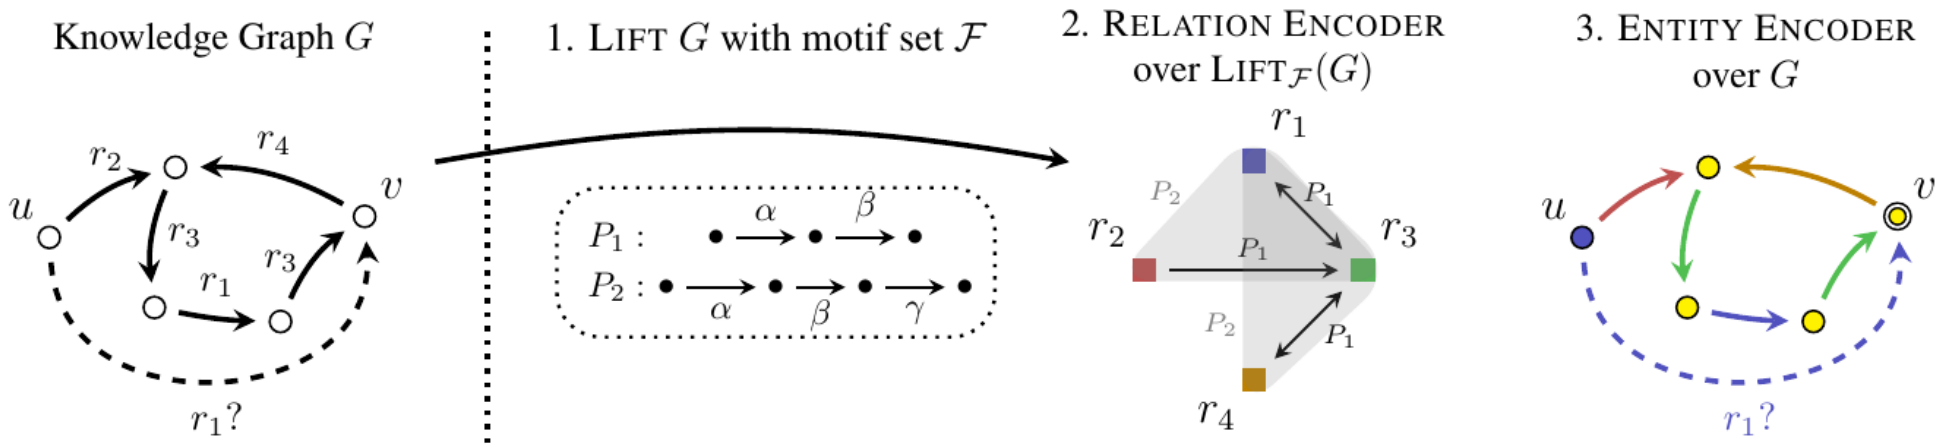
Condition interactions on the query.

2 Entity encoder

Start at the head; rank candidates.

MOTIF: Higher-Order Patterns

From pairwise relation interactions to richer motifs



Pairwise interactions

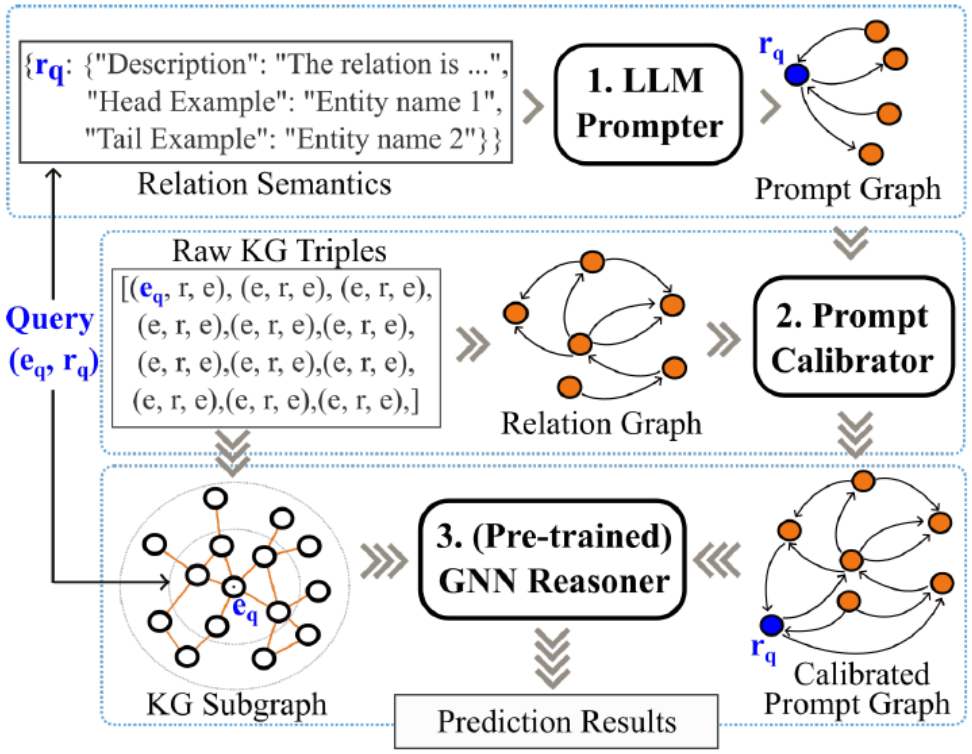
Encode interactions between two relation types.

Higher-order motifs

Encode multi-relation motifs in a relation hypergraph.

ProLINK: LLM-Generated Context

ProLINK adds language-generated context to graph structure



Three-stage pipeline

Generate

Create a prompt from relation meaning.

Adjust

Refine it with target-KG structure.

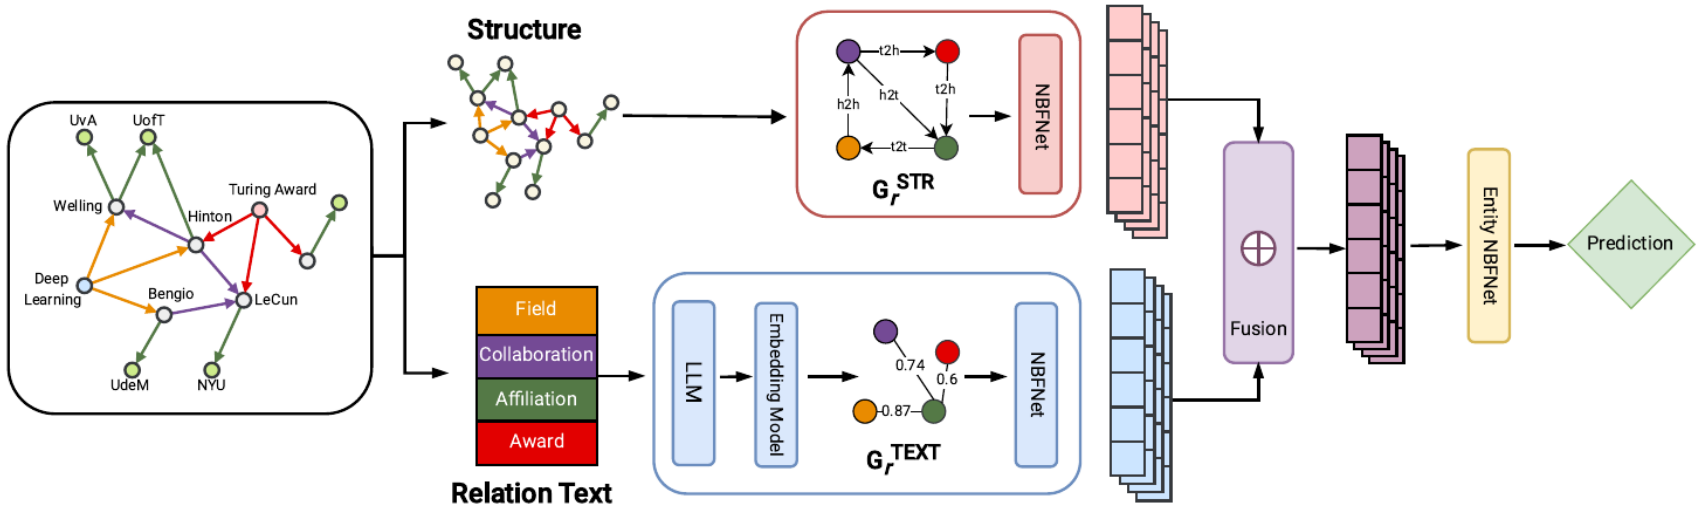
Reason

Run the pretrained graph reasoner.

Kai Wang, Yuwei Xu, Zhiyong Wu, Siqiang Luo. LLM as Prompter: Low-resource Inductive Reasoning on Arbitrary Knowledge Graphs. Findings of ACL 2024.

SEMMA: Structure and Semantics

SEMMA uses semantic information alongside relational structure



Structural view

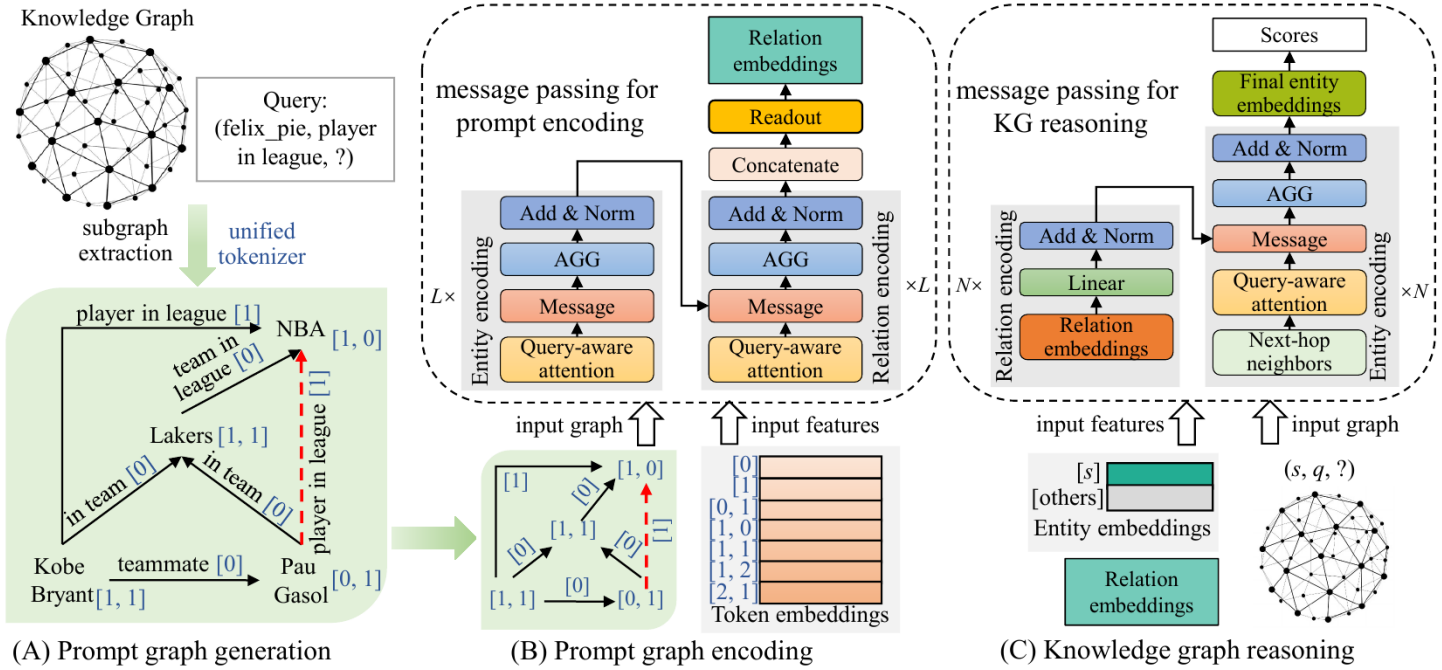
Relation interactions from KG structure.

Semantic view

A textual graph from relation descriptions.

KG-ICL: Learning from Examples

Can example facts provide relation context?



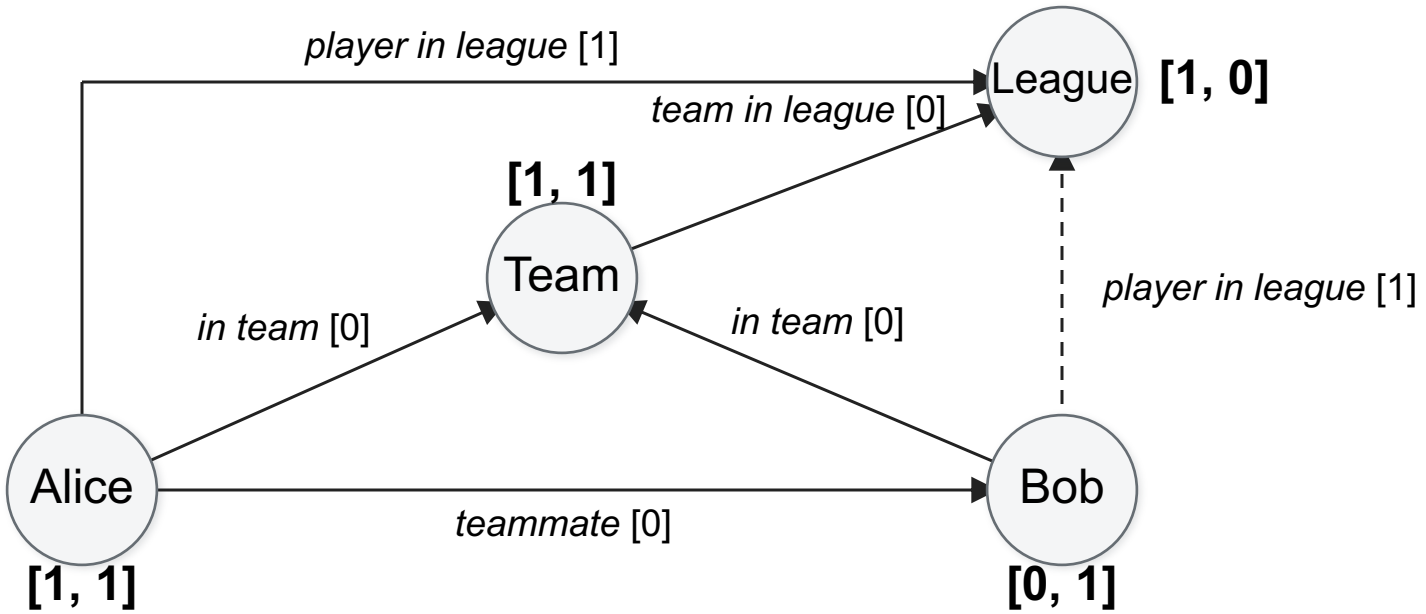
Known examples → **Prompt graphs** → **Relation context** → **KG reasoning**

Yuanning Cui, Zequn Sun, Wei Hu. A Prompt-Based Knowledge Graph Foundation Model for Universal In-Context Reasoning. NeurIPS 2024.

KG-ICL: Unified Tokenization

Encode entities and relations relative to an example fact

Example fact (Bob, *player in league*, League)



Simplified example; dashed edge: observed example fact

Entity token

[distance to head, distance to tail]
Team → [1, 1]

Relation token

Query relation → 1
Other relations → 0

KG-ICL: Prompt Graphs

From tokenized examples to relation prompts

1 Input

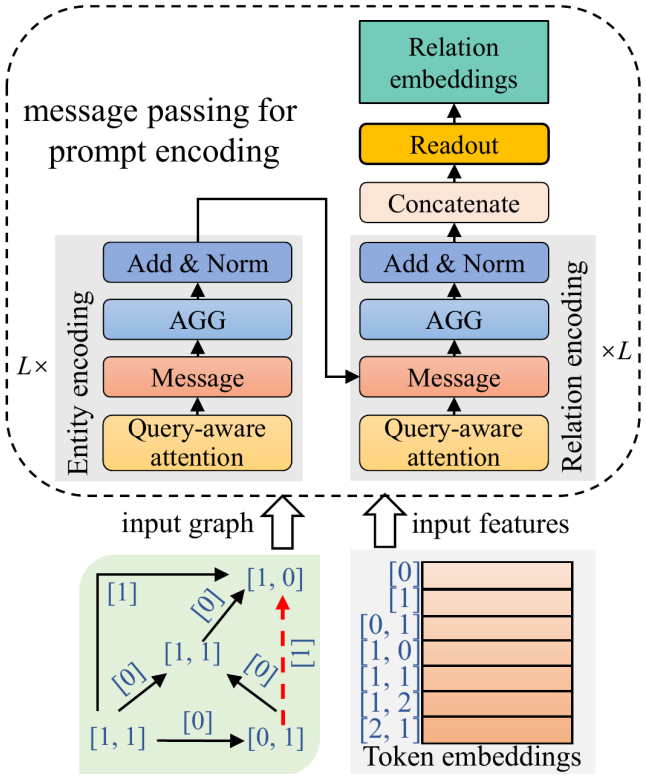
Initialize entity and relation tokens in each prompt graph.

2 Encode

Pass query-aware messages within each example graph.

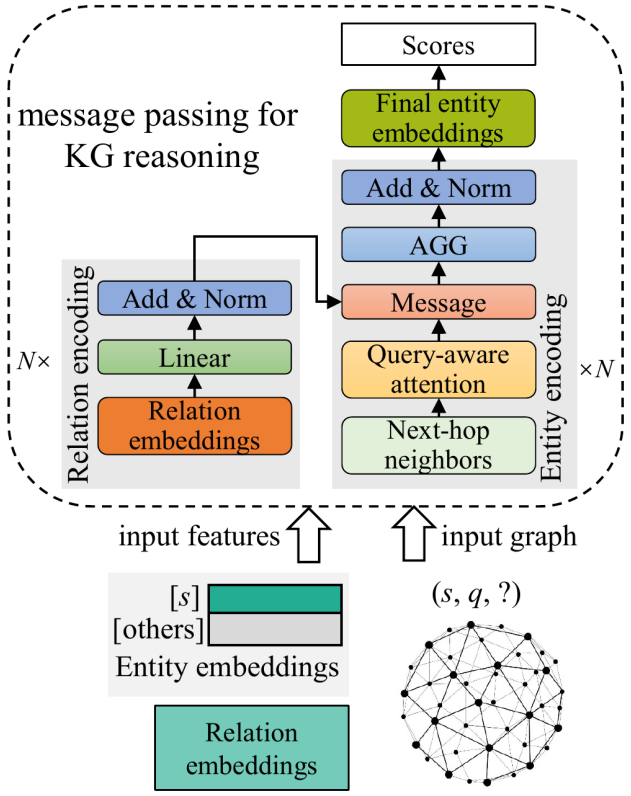
3 Readout

Combine relation features across examples to build the relation prompt.



KG-ICL: Prompt-Conditioned Reasoning

Prompt-guided message passing and candidate scoring



1 Relation context

Initialize relations with the prompt features.

2 Query head

Start entity-level propagation from the query head.

3 Prediction

Score and rank candidate tail entities.

Yuanning Cui, Zequn Sun, Wei Hu. A Prompt-Based Knowledge Graph Foundation Model for Universal In-Context Reasoning. NeurIPS 2024.

Reasoning: Transfer Evaluation

Same pretraining mixture, evaluated across 43 KGs

Table 5: Performance w.r.t the same pre-training mixture as ULTRA [6].

Models	Inductive (14 KGs)		Fully-inductive (13 KGs)		Transductive (16 KGs)		Average (43 KGs)	
	MRR	H@10	MRR	H@10	MRR	H@10	MRR	H@10
ULTRA (FB15k-237 WN18RR CoDEx-medium)	0.513	0.664	0.352	0.536	0.329	0.479	0.396	0.557
KG-ICL (FB15k-237 WN18RR CoDEx-medium)	0.547	0.700	0.431	0.629	0.357	0.506	0.441	0.606
KG-ICL (FB V1 NELL V1 CoDEx-small)	0.554	0.707	0.439	0.635	0.346	0.493	0.442	0.606

Average MRR

ULTRA → KG-ICL

0.396 → 0.441 (+0.045)

Fully inductive

0.352 → 0.431 MRR

Across 13 KGs

Target context

Target KG in both.

KG-ICL adds examples.

Yuanning Cui, Zequn Sun, Wei Hu. A Prompt-Based Knowledge Graph Foundation Model for Universal In-Context Reasoning. NeurIPS 2024. Appendix, Table 5.

Agenda

1. Introduction
2. Knowledge Graph Reasoning
3. **Knowledge Graph Entity Alignment**
4. Knowledge Graph Retrieval
5. Conclusion and Future Work

Entity Alignment

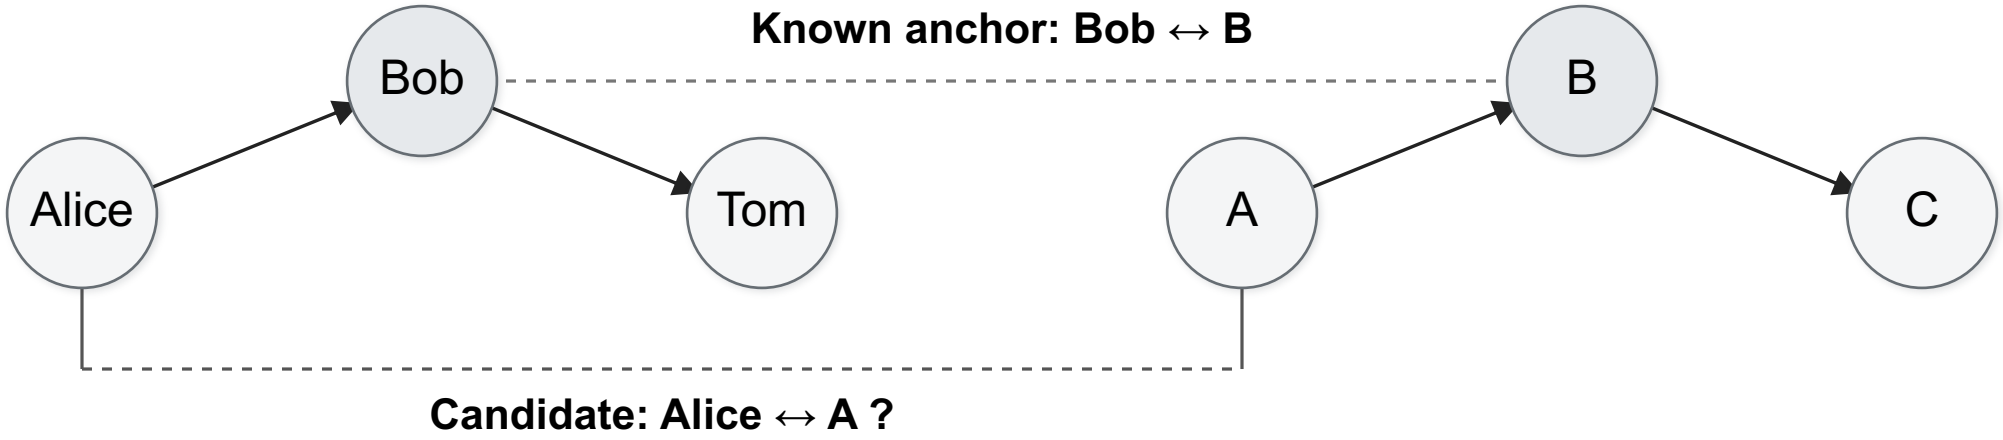
Identify the same object across different knowledge graphs

Source knowledge graph

Edges: *parentOf*

Target knowledge graph

Edges: *parentOf*



Input

Two graphs and known anchor pairs

Output

Ranked target entities

Entity Alignment: Method Landscape

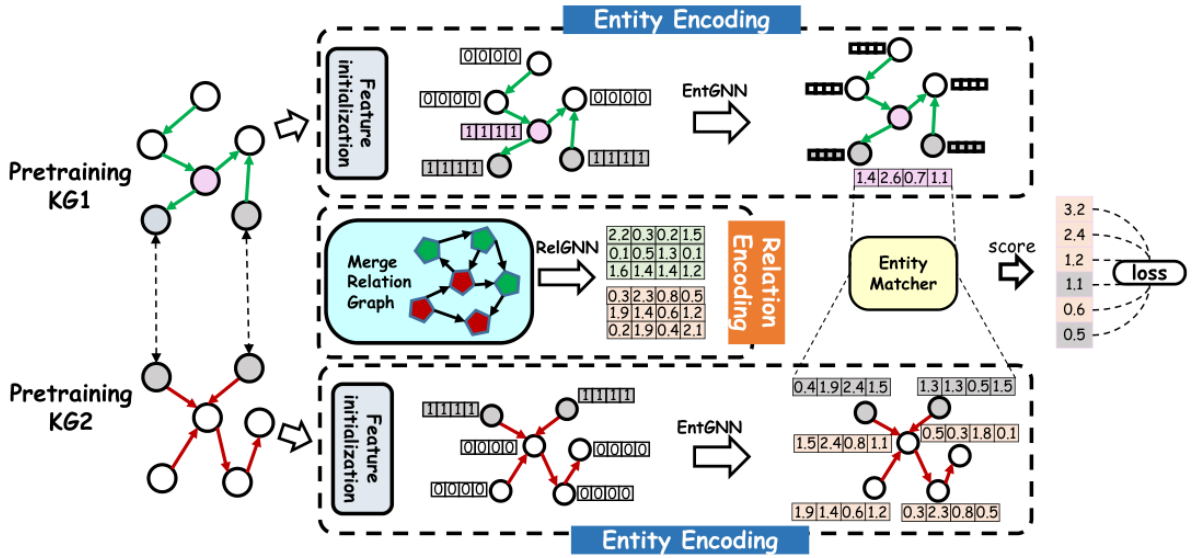
Alignment methods use different evidence and supervision

Approach	Evidence and mechanism	Typical setting
MTransE [1]	Pair alignments → mapped embeddings	Pair-specific
GCN-Align [2]	Neighborhood aggregation	Pair-specific
ZeroEA [3]	Language model + graph context	Zero-training
NeuSymEA [4]	Neural scores + symbolic rules	Protocol-specific
EAFM / ContextEA	Anchors + reusable graph model	Cross-graph transfer

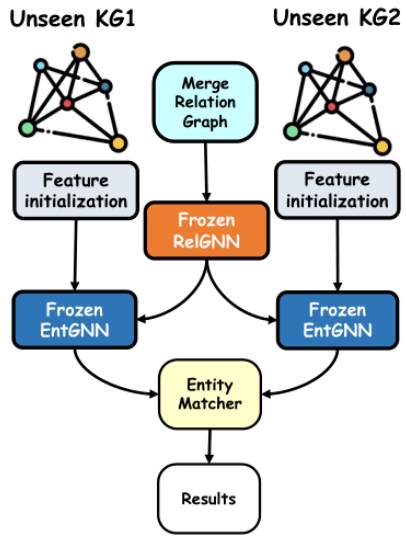
[1] Muhao Chen, Yingtao Tian, Mohan Yang, Carlo Zaniolo. Multilingual Knowledge Graph Embeddings for Cross-lingual Knowledge Alignment. IJCAI 2017.
[2] Zhichun Wang, Qingsong Lv, Xiaohan Lan, Yu Zhang. Cross-lingual Knowledge Graph Alignment via Graph Convolutional Networks. EMNLP 2018.
[3] Nan Huo, Reynold Cheng, Ben Kao, Wentao Ning, Nur Al Hasan Haldar, Xiaodong Li, Jinyang Li, Mohammad Matin Najafi, Tian Li, Ge Qu. ZeroEA: A Zero-Training Entity Alignment Framework via Pre-Trained Language Model. PVLDB 17(7), 2024.
[4] Shengyuan Chen, Zheng Yuan, Qinggang Zhang, Wen Hua, Jiannong Cao, Xiao Huang. NeuSymEA: Neuro-symbolic Entity Alignment via Variational Inference. NeurIPS 2025.
[5] Yuanning Cui, Zequn Sun, Wei Hu, Kexuan Xin, Zhangjie Fu. Breaking the Reasoning Horizon in Entity Alignment Foundation Models. IJCAI 2026.
[6] Xingyu Chen, Yuanning Cui, Zequn Sun, Wei Hu. Harnessing Structural Context for Entity Alignment Foundation Models. arXiv:2606.06109, 4 June 2026 (preprint).

EAFM: Pretraining and Transfer

Source pretraining and target inference



(a) Entity alignment pre-training



(b) Transferable inference

Source pretraining

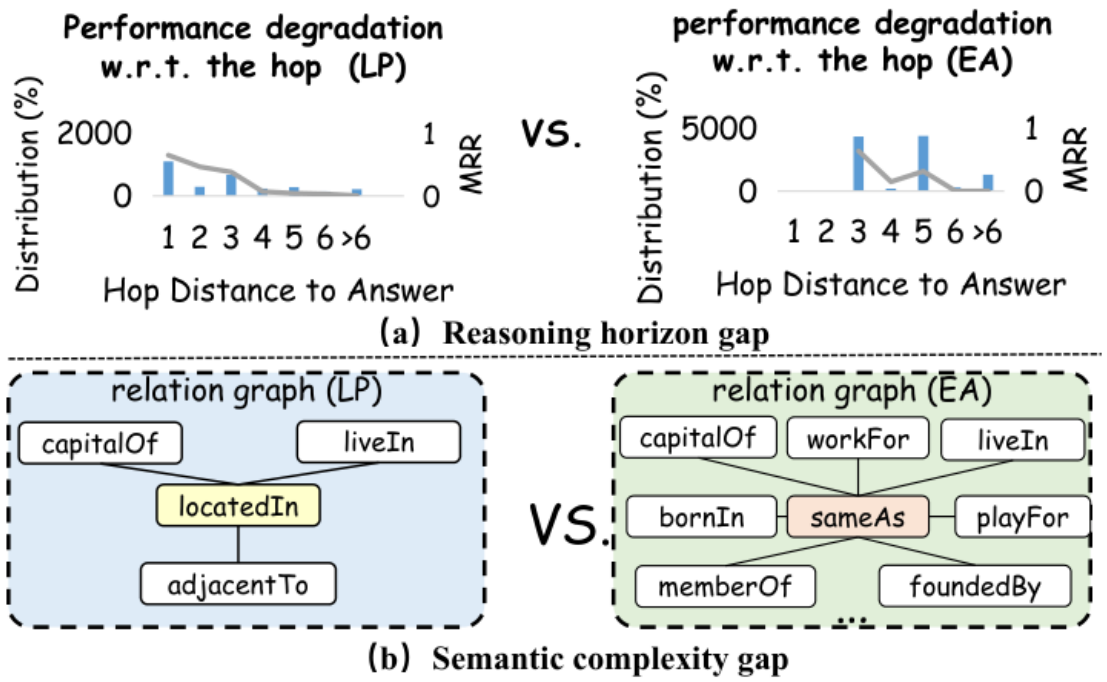
Train the encoder and matcher on KG pairs with seed alignments.

Target inference

Freeze parameters; use target KGs and available anchors.

Entity Alignment: Sparse Anchors

EAFM: The Reasoning Horizon Gap



Original figure: distribution-axis units are unclear.

1 Reasoning horizon

Sparse anchors can leave matches beyond the message-passing range.

2 Semantic complexity

One *sameAs* relation covers many entity types and patterns.

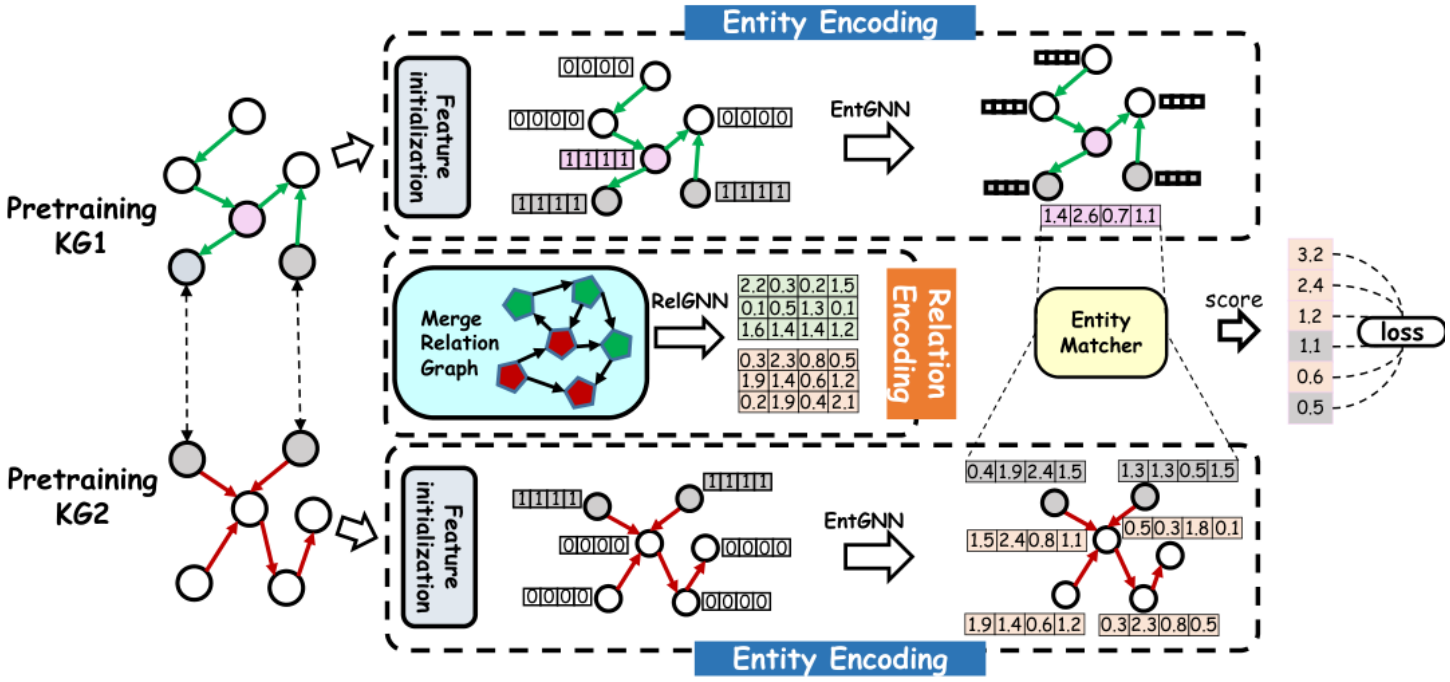
Design response

Use local anchors and query-specific relation context.

Yuanning Cui, Zequn Sun, Wei Hu, Kexuan Xin, Zhangjie Fu. Breaking the Reasoning Horizon in Entity Alignment Foundation Models. IJCAI 2026.

EAFM: Anchor-Aware Encoding

Parallel anchor-aware encoding across two KGs



(a) Entity alignment pre-training

1 Initialize

Use query-local anchor pairs

2 Encode

Share parameters across both graphs

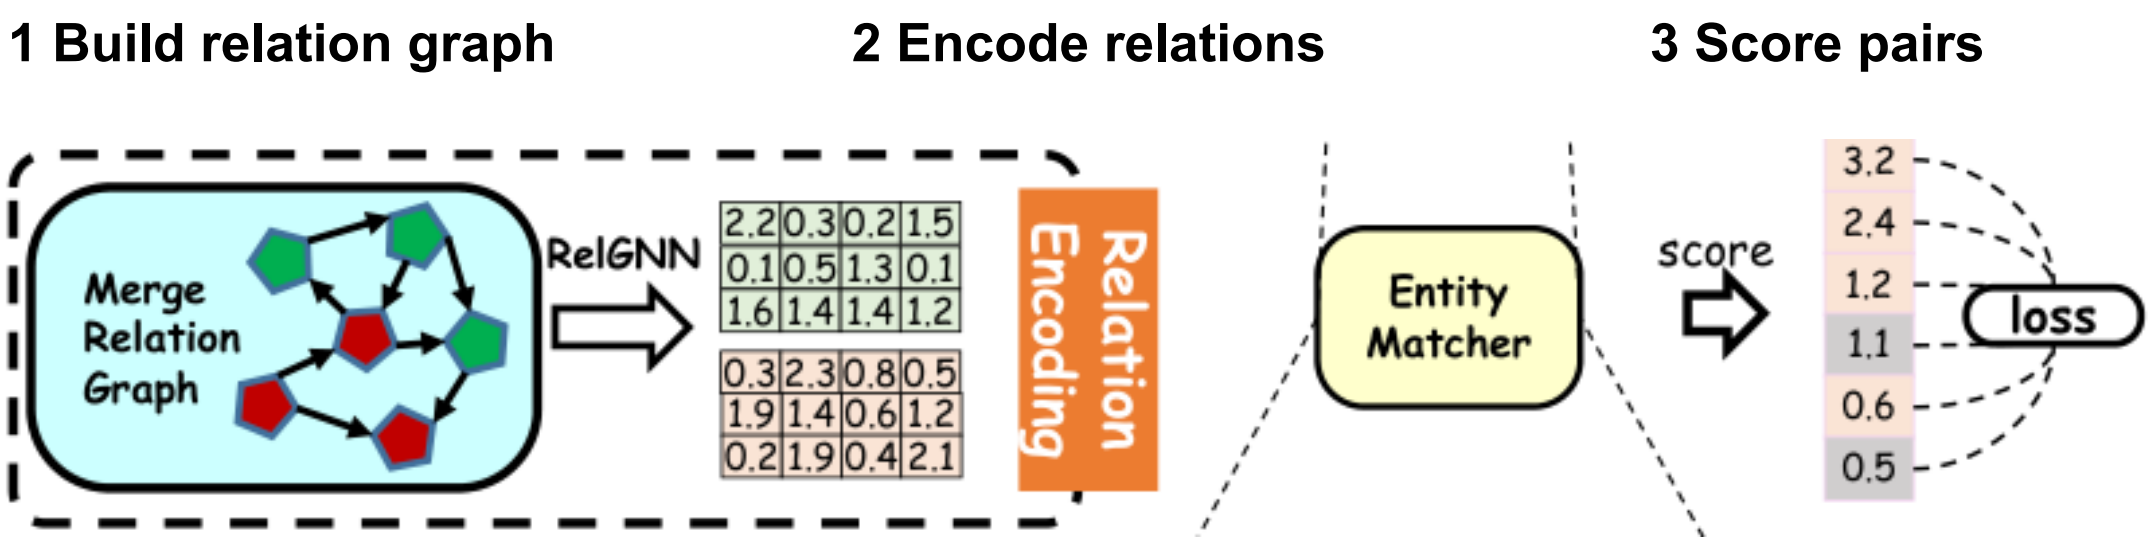
3 Match

Compare relative entity contexts

Yuanning Cui, Zequn Sun, Wei Hu, Kexuan Xin, Zhangjie Fu. Breaking the Reasoning Horizon in Entity Alignment Foundation Models. IJCAI 2026.

EAFM: Anchor-Relative Context

Relation context supports pairwise matching



Anchor-induced relation interactions Query-conditioned RelGNN Learnable matcher

Yuanning Cui, Zequn Sun, Wei Hu, Kexuan Xin, Zhangjie Fu. Breaking the Reasoning Horizon in Entity Alignment Foundation Models. IJCAI 2026.

ContextEA: Cross-Graph Context

From parallel encoding to cross-graph interaction

Design choice	EAFM	ContextEA
Entity encoding	Within each KG	Across anchor bridges
Relation context	Both use a query-conditioned relation graph	
Scoring	Interaction matcher	Coarse score + correction
Target input	Both use graph structure and anchors	

[1] Yuanning Cui, Zequn Sun, Wei Hu, Kexuan Xin, Zhangjie Fu. Breaking the Reasoning Horizon in Entity Alignment Foundation Models. IJCAI 2026.
[2] Xingyu Chen, Yuanning Cui, Zequn Sun, Wei Hu. Harnessing Structural Context for Entity Alignment Foundation Models. arXiv:2606.06109, 4 June 2026 (preprint).

ContextEA: Anchor Bridges

Anchor bridges enable cross-graph message passing

Connect known anchors

Connect each known anchor pair.
The two entities remain distinct.

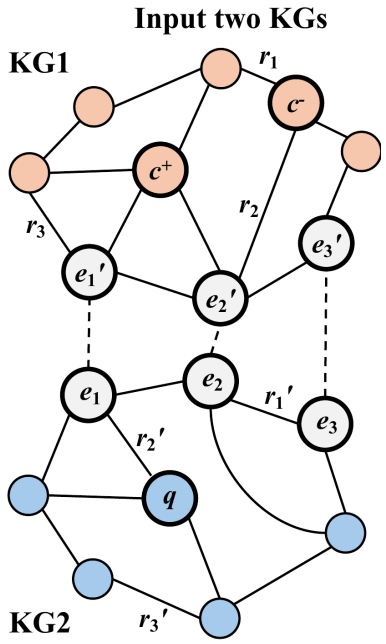
Propagate across KGs

Pass messages along KG edges
and anchor bridges.

Build candidate representations

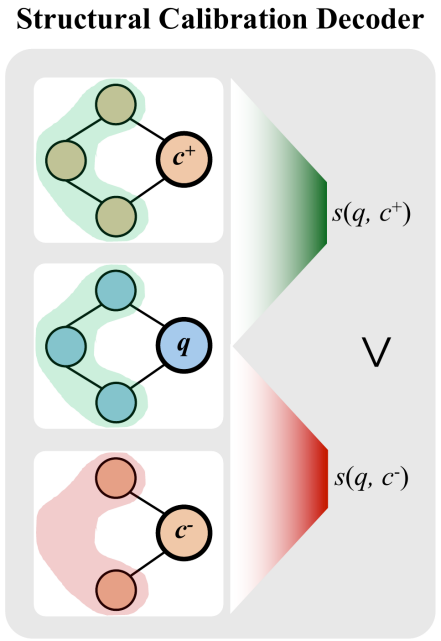
Use structural evidence from both KGs.

q: query · c^+ / c^-: candidates



ContextEA: Candidate Calibration

ContextEA: Structural Calibration



Coarse ranking

Rank eligible target entities; select the top-k.

Structural calibration

Refine candidates using entity, neighborhood, relation and anchor evidence.

Final score

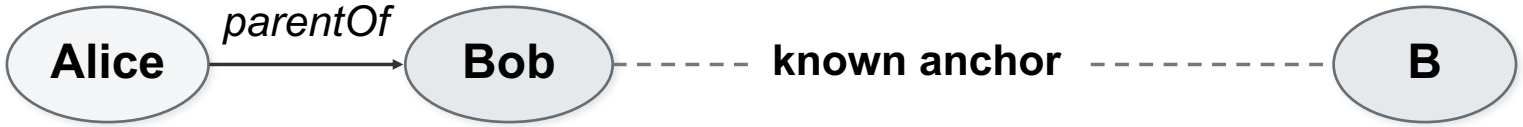
Coarse score + weighted structural correction

Outside top-k: keep the coarse score.

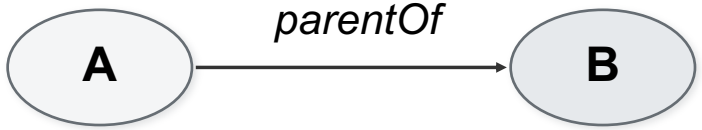
ContextEA: A Matching Example

A teaching example of candidate calibration

Query in the source graph

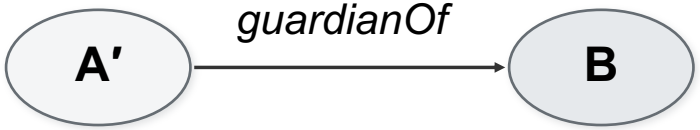


Candidate A



Matching relation and anchor support A.

Candidate A′



Different relation: check its semantics.

Entity Alignment: Transfer Results

Frozen transfer and fine-tuning across 29 datasets

Method	OpenEA			SRPRS			DBP			Average		
	MRR	H@10	H@1	MRR	H@10	H@1	MRR	H@10	H@1	MRR	H@10	H@1
ULTRA-EA pretrain	0.477	0.635	0.385	0.350	0.557	0.241	0.203	0.379	0.107	0.395	0.569	0.297
ULTRA-EA finetune	0.509	0.691	0.416	0.380	0.587	0.271	0.223	0.409	0.137	0.424	0.614	0.328
EAFM pretrain	0.602	0.764	0.517	0.530	0.735	0.426	0.294	0.452	0.214	0.529	0.702	0.440
EAFM finetune	0.656	0.814	0.573	0.605	0.792	0.508	0.464	0.662	0.394	0.609	0.782	0.524
ContextEA pretrain	0.713	0.857	0.636	0.675	0.837	0.587	0.522	0.667	0.444	0.669	0.819	0.589
ContextEA finetune	**0.722**	**0.859**	**0.648**	**0.680**	**0.839**	**0.594**	**0.588**	**0.743**	**0.503**	**0.688**	**0.834**	**0.608**

Average MRR · ULTRA-EA → EAFM → ContextEA

Pretrain 0.395 → 0.529 → 0.669

Fine-tune 0.424 → 0.609 → 0.688

Target-side information

Target graphs and known anchors are used in both settings.

Xingyu Chen, Yuanning Cui, Zequn Sun, Wei Hu. Harnessing Structural Context for Entity Alignment Foundation Models. arXiv:2606.06109, 4 June 2026 (preprint). Table 1.

Agenda

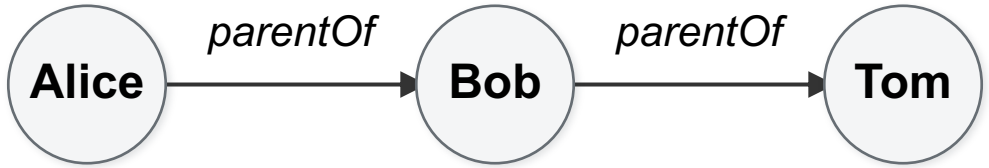
1. Introduction
2. Knowledge Graph Reasoning
3. Knowledge Graph Entity Alignment
4. **Knowledge Graph Retrieval**
5. Conclusion and Future Work

Question Answering and KG Evidence

Multi-hop KGQA benefits from structured evidence

Question: Who is Tom's grandparent?

Observed graph facts



Retrieve the connected two-hop path

Example rule (assumed)

$parentOf(x, y) \land parentOf(y, z) \Rightarrow grandparentOf(x, z)$

Derived fact

Alice is Tom's grandparent

Graph Retrieval: Four Approaches

Four approaches differ in output and transfer

System	Input	Retrieved unit	Model type
ToG [1]	Existing KG	Paths	LLM search
G-Retriever [2]	Textual graph	Connected subgraph	Subgraph search
GFM-RAG [3]	Documents + KG index	Ranked documents	Pretrained GFM
KGFR [4]	KG + question	Nodes / edges / paths	Pretrained GFM

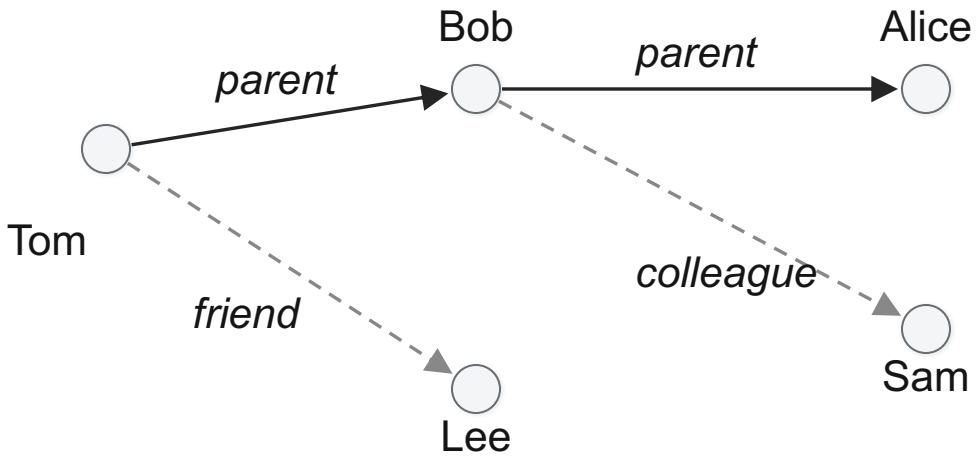
[1] Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel M. Ni, Heung-Yeung Shum, Jian Guo. Think-on-Graph: Deep and Responsible Reasoning of Large Language Model on Knowledge Graph. ICLR 2024.
[2] Xiaoxin He, Yijun Tian, Yifei Sun, Nitesh V. Chawla, Thomas Laurent, Yann LeCun, Xavier Bresson, Bryan Hooi. G-Retriever: Retrieval-Augmented Generation for Textual Graph Understanding and Question Answering. NeurIPS 2024.
[3] Linhao Luo, Zicheng Zhao, Reza Haffari, Dinh Phung, Chen Gong, Shirui Pan. GFM-RAG: Graph Foundation Model for Retrieval Augmented Generation. NeurIPS 2025.
[4] Yuanning Cui, Zequn Sun, Wei Hu, Zhangjie Fu. KGFR: A Foundation Retriever for Generalized Knowledge Graph Question Answering. IEEE Transactions on Knowledge and Data Engineering (TKDE), 2026.

Think-on-Graph: LLM-Guided Search

Think-on-Graph: LLM-guided path exploration

Illustrative question

Who is Tom's grandparent?



Follow *parentOf* edges in reverse.

How the search works

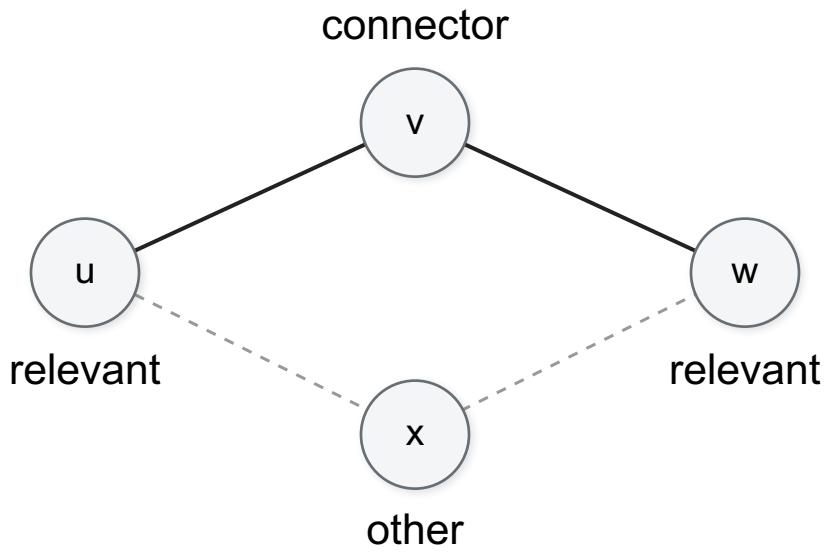
1. Explore nearby relations and entities.
2. Ask the LLM to score and select branches.
3. Repeat across hops; answer from the paths.

Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel M. Ni, Heung-Yeung Shum, Jian Guo. Think-on-Graph: Deep and Responsible Reasoning of Large Language Model on Knowledge Graph. ICLR 2024.

G-Retriever: Subgraph Retrieval

G-Retriever: evidence as a connected subgraph

Select connected evidence



Solid: selected · Dashed: not selected

Represent and answer

1 Score relevance

Score node and edge relevance.

2 Select connected evidence

PCST: relevance and connection cost.

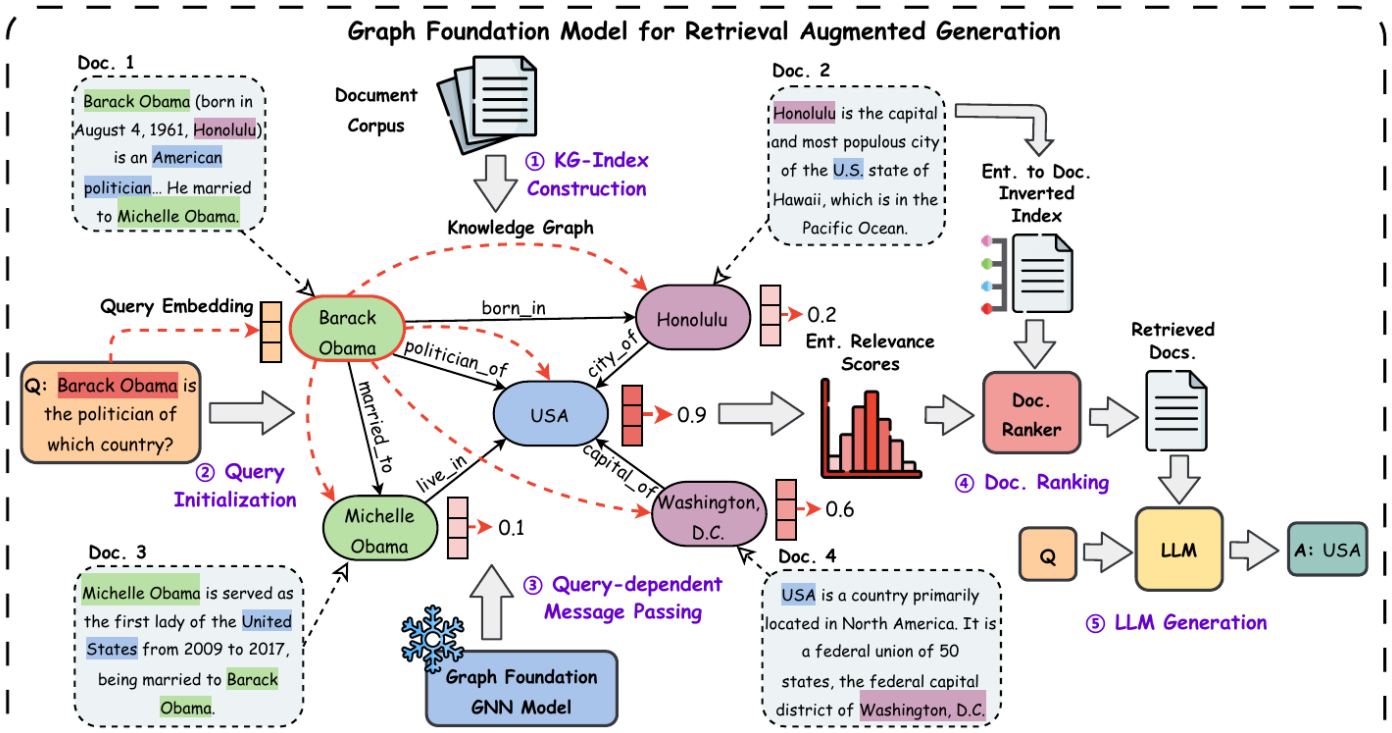
3 Encode and answer

Guide the LLM with the subgraph.

Xiaoxin He, Yijun Tian, Yifei Sun, Nitesh V. Chawla, Thomas Laurent, Yann LeCun, Xavier Bresson, Bryan Hooi. G-Retriever: Retrieval-Augmented Generation for Textual Graph Understanding and Question Answering. NeurIPS 2024.

GFM-RAG: Document Retrieval

KG completion pretraining → document retrieval training



1 KG index

Documents → KG and inverted index

2 Entity scoring

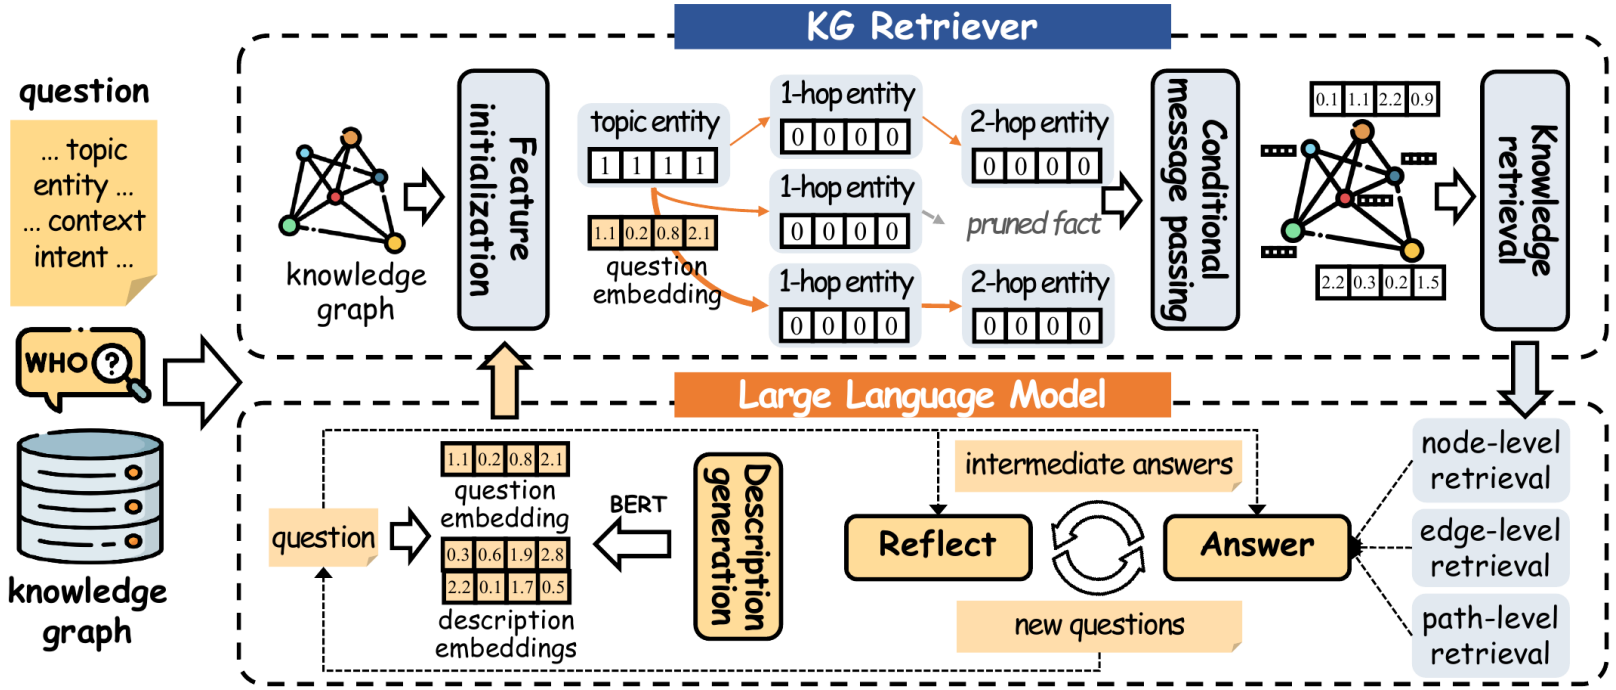
Query-guided GFM scores graph entities

3 Document ranking

Map entity scores back to documents

KGFR: Transferable Retrieval

Learn from QA pairs; reuse the retriever on new graphs



1 KG retriever

Score entities; attend to edges

2 LLM

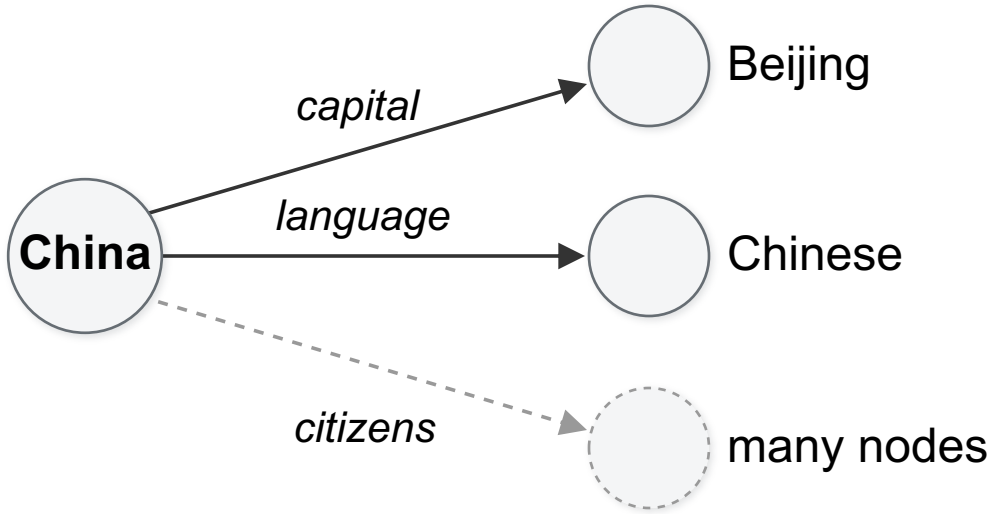
Inspect evidence; reflect and retrieve again if needed

Yuanning Cui, Zequn Sun, Wei Hu, Zhangjie Fu. KGFR: A Foundation Retriever for Generalized Knowledge Graph Question Answering. IEEE Transactions on Knowledge and Data Engineering (TKDE), 2026.

KGFR: Progressive Propagation

Control expansion without discarding visited evidence

A high-degree entity



Dashed branch: suppress new expansion

Asymmetric progressive propagation

Limit expansion for an oversized (entity, relation) group.

Keep the entity and expand other useful relations.

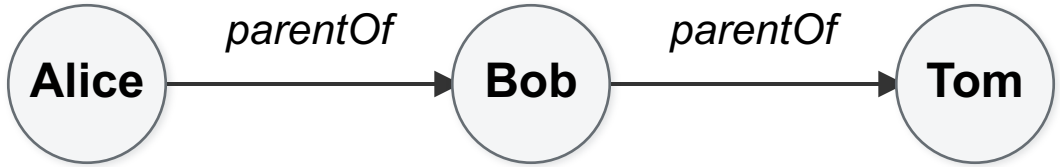
Preserve connections to previously visited nodes.

KGFR: Multi-Level Retrieval

Three evidence levels support a retrieval–reflection loop

A shared example graph

Illustrative



Need more evidence?

Reflect, refine and retrieve again.

Different levels of evidence

Node

Alice — a candidate entity

Edge

Alice → Bob — a supporting fact

Path

Alice → Bob → Tom

Connected facts across multiple hops

Yuanning Cui, Zequn Sun, Wei Hu, Zhangjie Fu. KGFR: A Foundation Retriever for Generalized Knowledge Graph Question Answering. IEEE Transactions on Knowledge and Data Engineering (TKDE), 2026.

Answer Quality and Efficiency

Cross-dataset transfer and inference cost

Cross-dataset transfer · Qwen-max

TABLE IV: Generalization between WebQSP and CWQ

LLMs	Training → Testing	F1	Hit	H@1
Qwen-max	CWQ → WebQSP	72.9	86.4	81.8
	WebQSP → WebQSP	74.7	90.3	83.2
	WebQSP → CWQ	57.6	70.2	62.1
	CWQ → CWQ	61.6	71.8	63.6

Efficiency · per question

TABLE IX: Efficiency comparison with baseline methods

Datasets	Methods	LLM calls	Tokens	Time (s)
CWQ	ToG	22.6	9,669.4	96.5
	PoG	13.3	8,156.2	23.3
	LLM-KGFR	2.4	3,145.1	9.2
WebQSP	ToG	15.9	7,018.9	63.1
	PoG	9.0	5,517.7	16.8
	LLM-KGFR	2.1	2,725.2	7.6

F1 on CWQ

Train WebQSP: 57.6 · Train CWQ: 61.6

CWQ inference cost

2.4 LLM calls · 9.2 s

Shared Freebase KG: cross-dataset transfer; efficiency evaluated separately.

Yuanning Cui, Zequn Sun, Wei Hu, Zhangjie Fu. KGFR: A Foundation Retriever for Generalized Knowledge Graph Question Answering. IEEE Transactions on Knowledge and Data Engineering (TKDE), 2026. Tables IV and IX.

Agenda

1. Introduction
2. Knowledge Graph Reasoning
3. Knowledge Graph Entity Alignment
4. Knowledge Graph Retrieval
5. **Conclusion and Future Work**

Summary

KG foundation models reuse learned patterns across graphs

Reasoning	Predict missing facts with transferable relation context.
Entity alignment	Match entities using anchors and cross-graph context.
Retrieval	Retrieve useful evidence to support question answering.

Open Challenges

Open challenges for transferable KG foundation models

Generalization	Scalability	Collaboration
Schemas and relation meanings can shift across graphs.	High-degree expansions increase retrieval cost.	The LLM may misread or overtrust graph evidence.
Combine transferable structure with semantic context.	Expand selectively; measure evidence recall.	Show evidence, verify support and reflect selectively.

Thank You!

Zequn Sun

Nanjing University, China