

Pengyu Zhang

PhD candidate in Computer Science · University of Amsterdam

Amsterdam, the Netherlands · +31 6 16144184 · p.zhang@uva.nl

pengyu-zhang.github.io · github.com/pengyu-zhang · [LinkedIn](#) · [Google Scholar](#) · [ORCID](#)

• ABOUT ME

I work on machine learning for knowledge graphs, where the data keeps changing and many entities look alike. I build models that use graph evolution, user activity and time itself to keep entity linking, knowledge graph completion and recommendation accurate on sparse, long-tail and look-alike entities. My first-author methods improve on the prior state of the art by up to **20.9%** in entity linking three years after training, **9.8%** in recommendation, and **58.2%** on the hardest look-alike entities. Code and benchmarks for the published ones are on GitHub.

Open to Data Scientist and Machine Learning Engineer roles, and to postdoctoral positions.

Multi-modal knowledge graphs

Entity linking

Link prediction

Recommendation

Large language models

• EXPERIENCE

2022 – present **PhD Candidate**, Intelligent Data Engineering Lab (INDElab), University of Amsterdam

- Lead five first-author projects on temporally robust multi-modal learning, from problem framing to publication (CIKM 2024, ECAI 2024, ESWC 2026, two under review), and contribute to collaborations published at ISWC, ACL, LREC, EKAW and in Expert Systems with Applications.
- Release code and data for every published first-author project: two yearly-snapshot entity linking benchmarks (Graph-TempEL, GCL-TempEL) and three enriched multi-modal knowledge graph datasets.
- Work at scale: graphs of up to 5.4M entities and 2M+ training samples; a pipeline that captioned more than 300,000 Wikipedia images with three vision-language models and fused the captions with LLMs; entity linking over 21 million tweets; a deployed visual analytics system that research administrators used on 4,000 real papers.

• RESEARCH PROJECTS

TimeRoute Time-aware multi-modal recommendation under review, arXiv 2026

- Which modality matters for a purchase shifts over time, and differently for each user, yet recommenders fuse text, images and audio with one global weighting. I built a per-user router that turns interaction timestamps into modality weights, and a time-conditioned diffusion module that filters outdated edges before graph propagation.
- Best on every metric on TikTok, Amazon-Baby and Amazon-Sports over ten seeds, **4.9% to 9.8%** above the strongest baselines. Swap the timing signal for noise and the gain vanishes.

Time Imprint Time as a modality for look-alike entities under review, arXiv 2026

- Some entities share almost identical text and images, like Napoleon Bonaparte and the 2023 film Napoleon. I treat time as an entity-level modality: select a few timestamps around the median, pool them with attention, and inject the result into the encoder, the scoring function and a contrastive objective.
- New state of the art on MKG-W with Hits@1 up **6.1%**. On the most ambiguous 1% of entities, Hits@1 improves by up to **58.2%**.

Select, Don't Train Modular entity disambiguation with LLM selection ISWC 2026

- Dual-encoders squeeze candidate retrieval and fine-grained selection into one embedding space, and need retraining whenever the knowledge base changes. With Fina Polat (lead), we split the two: a retriever returns 16 candidates and an LLM picks one or abstains.
- Plain BM25 with no retriever training reaches **86.3 micro-F1** on ZELDA, 4 points above the best trained dual-encoder; a trained dense retriever pushes it to 88.5.

Beyond Images Data-centric enrichment of multi-modal knowledge graphs

ESWC 2026

- Curated image sets drop logos and symbols because visual encoders turn them into noise. I built a pipeline that retrieves every Wikipedia image for an entity, captions each one with a vision-language model, and fuses the captions with an LLM into one text summary, so the downstream model does not change.
- Hits@1 rises **4.8% to 7.8%** across three datasets and four models; on a hand-picked set of 20 logo and symbol entities it more than quadruples. Three enriched datasets released.

Fusion Training Data mixing for hybrid-reasoning LLMs

ACL 2026 SRW

- Hybrid models like Qwen3 answer quickly or think step by step from one set of weights, but the two kinds of training data interfere. With Congfeng Cao (lead), we swept seven data ratios and three training schedules on Qwen3-4B and released the Fusion Bench benchmark.
- More short-answer data steadily erodes long-form reasoning, the two accuracies are negatively correlated, and interleaving the two is the most robust schedule.

Graph - TempCZ Large-scale temporal link prediction

LREC 2026

- Software-mention data had no graph structure and no time axis. With Congfeng Cao (lead), we turned mentions from 1.5 million biomedical papers into a publication-software graph with more than 6.5 million edges spanning 1959 to 2022, and framed software usage as temporal link prediction.
- GraphSAGE reaches **92.9%** accuracy, 6 points above the best feature-based baseline. Accuracy decays as the gap between training and test years grows, with a transfer window of about two years.

CYCLE & TIGER Temporally robust entity linking

CIKM 2024 & ECAI 2024

- An entity linker is trained on one snapshot of the knowledge graph and used for years while the graph moves on, so accuracy decays. I designed two fixes: TIGER adds graph structure to a text-only bi-encoder so entity representations stay separable as descriptions drift; CYCLE treats the change itself as supervision, turning relations that appear or vanish between yearly snapshots into contrastive positives and negatives.
- Against the 2023 state of the art, TIGER gains **16.2%** at a one-year gap and **20.9%** at three years on continual entities; CYCLE gains **13.9%** and **17.8%** averaged over continual and new entities, most on low-degree ones. Both ship with public benchmarks (Graph-TempEL, GCL-TempEL).

• EDUCATION

2022 – present	PhD in Computer Science , University of Amsterdam, INDElab. Defense expected in early 2027. Thesis: <i>Addressing Temporal Dynamics and Entity Ambiguity in Knowledge Graph Tasks with Multi-Source Information</i> . Promotor Prof. Dr. Paul Groth ; co-promotor Dr. Klim Zaporozets . Funded by the China Scholarship Council.
2019 – 2022	MEng in Control Engineering , Beijing University of Technology. Supervisor Prof. Dr. Yong Zhang . Outstanding Master's Thesis award.

• SKILLS

Programming	Python, SQL (PostgreSQL), Bash, Git, Linux.
ML & DL	PyTorch, Hugging Face, scikit-learn, NumPy, Pandas, Weights & Biases; multi-GPU training on A100 and H100; graph neural networks, contrastive learning, multi-modal fusion, diffusion models, temporal representation learning, LLM and vision-language pipelines.
Languages	Chinese (native), English (fluent, IELTS 7.0), Dutch (A2).

• AWARDS & SERVICE

Awards	Conference travel grant, ACL 2026. Outstanding Master's Thesis, Beijing University of Technology, 2022.
Talks	Paper presentations at ACL 2026 (San Diego), ESWC 2026 (Dubrovnik), LREC 2026 (Palma), EKAW 2024 (Amsterdam), CIKM 2024 (Boise) and ECAI 2024 (Santiago de Compostela).
Service	Reviewer or sub-reviewer for CIKM, ECAI, ESWC and PKDD.
Teaching	Teaching assistant, Data Engineering, Beijing University of Technology, 2020 – 2021.

• PUBLICATIONS

- [14] TimeRoute: Time-Aware Modality Routing and Diffusion for Multi-Modal Recommendation. **Pengyu Zhang**, Yangqin Jiang, Klim Zaporozets, Congfeng Cao, Paul Groth. Under review. [arXiv](#)
- [13] Time Imprint: Learning Time-Aware Representations in Multi-Modal Knowledge Graphs. **Pengyu Zhang**, Klim Zaporozets, Congfeng Cao, Jia-Hong Huang, Paul Groth. Under review. [arXiv](#)
- [12] Select, Don't Train: The Benefits of Modular Entity Disambiguation with LLM-Based Selection. Fina Polat, Daniel Daza, **Pengyu Zhang**, Klim Zaporozets, Paul Groth. ISWC 2026. [Code](#)
- [11] Fusion Training for Mathematical Generalization in Large Language Models. Congfeng Cao, **Pengyu Zhang**, Jelke Bloem. ACL 2026, Student Research Workshop. [Paper](#) · [DOI](#) · [Code](#)
- [10] Are a Thousand Words Better Than a Single Picture? Beyond Images – A Framework for Multi-Modal Knowledge Graph Dataset Enrichment. **Pengyu Zhang**, Klim Zaporozets, Jie Liu, Jia-Hong Huang, Paul Groth. ESWC 2026. [Paper](#) · [DOI](#) · [Code](#) · [Demo](#) · [YouTube](#) · [Bilibili](#)
- [9] Graph-TempCZ: A Graph Representation of Software Mentions for Predicting Software Usage in Scientific Publications. Congfeng Cao, **Pengyu Zhang**, Jelke Bloem. LREC 2026. [Paper](#) · [DOI](#) · [Code](#)
- [8] A Survey of Large Language Models for Data Challenges in Graphs. Mengran Li, **Pengyu Zhang**, Wenbin Xing, Yijia Zheng, Klim Zaporozets, Junzhou Chen, Ronghui Zhang, Yong Zhang, Siyuan Gong, Jia Hu, Xiaolei Ma, Zhiyuan Liu, Paul Groth, Marcel Worring. Expert Systems with Applications 2026. [Paper](#) · [DOI](#) · [Code](#)
- [7] Understanding the Impact of Entity Linking on the Topology of Entity Co-occurrence Networks for Social Media Analysis. James Nevin, **Pengyu Zhang**, Dimitar Dimitrov, Michael Lees, Paul Groth, Stefan Dietze. EKAW 2024. [Paper](#) · [DOI](#) · [Code](#)
- [6] CYCLE: Cross-Year Contrastive Learning in Entity-Linking. **Pengyu Zhang**, Congfeng Cao, Klim Zaporozets, Paul Groth. CIKM 2024. [Paper](#) · [DOI](#) · [Code](#)
- [5] TIGER: Temporally Improved Graph Entity Linker. **Pengyu Zhang**, Congfeng Cao, Paul Groth. ECAI 2024. [Paper](#) · [DOI](#) · [Code](#)
- [4] Relationship Updating Network with Contrastive Learning. **Pengyu Zhang**, Yong Zhang, Xinglin Piao, Yongliang Sun, Baocai Yin. Physica A 2024. [Paper](#) · [DOI](#) · [Code](#)
- [3] MVMA-GCN: Multi-view Multi-layer Attention Graph Convolutional Networks. **Pengyu Zhang**, Yong Zhang, Jingcheng Wang, Baocai Yin. Engineering Applications of Artificial Intelligence 2023. [Paper](#) · [DOI](#) · [Code](#)
- [2] Visual Analysis for Name Disambiguation of Academic Papers (in Chinese). **Pengyu Zhang**, Yong Zhang, Yanjie Cui, Baocai Yin. Journal of Computer-Aided Design & Computer Graphics 2022. [Paper](#) · [DOI](#) · [Code](#) · [Demo](#) · [YouTube](#) · [Bilibili](#)
- [1] Dual-Channel Heterogeneous Graph Network for Author Name Disambiguation. Xin Zheng, **Pengyu Zhang**, Yanjie Cui, Rong Du, Yong Zhang. Information 2021. [Paper](#) · [DOI](#) · [Code](#)