

EXPERIENCE

Google, Mountain View

Nov 2021 - now - Research Scientist

- [MUM](#) foundation model Platform
 - A widely used foundation model platform, which has won many awards (Google Tech Impact 2023, Search Tech Impact 2022)
 - Evaluation and launch of various architectures such as [dual encoder](#), [lightweight decoder](#), [MoE](#). Unblocked dual encoder launch by solving cross-lingual problems. Developed video dual encoder and launched semantic video segment search for Google Video Search
 - Spatial aware architecture for Google street view models.
 - Developed grounding models/services which are used by 10+ clients
 - Prototyped self-improvement post training for LLM reasoning
 - Prototyped Agent Platform to demo multi-step travel agent concept for [SGE](#). Applied distillation and RLHF methods for Map/Search agent, which was [demoed at GoogleIO](#) and approved for launch.
- Research ([google scholar](#))
 - Senior area chair for ACL, Neurips and EMNLP etc
 - Many publications in spatial aware model architectures.

Apple, Cupertino

May 2020 - Oct 2021 - Research Scientist

- KG answers and Web answers
 - Define and develop passage answer metrics, which are critical for product development and launch decisions. Improve short answer metric efficiency using parsing techniques such as Duckling.
 - Unify the evaluation of different answer domains with infra teams
 - Launched neural semantic parsing in search and siri.
- Research
 - Organize Apple's participation of the EfficientQA competition
 - Co-organize the NLP internship program, and co-host summer interns
 - Serve as one of the technical sponsors for Apple's NLU fundings

Mosaix.ai, Palo Alto

Feb 2018 - Apr 2020 - Co-founder and Chief Scientist

- Built and ran a team of NLP/ML engineers/researchers from the ground up. Managed the development and scaling of NLP/ML/quality infrastructures as the foundation for a voice AI platform, which serves millions of users.
- Research activities (collaborations, talks/interviews, reviews/chairs, workshop organizer) and publish at top AI conferences with innovations.

Google, Mountain View

Jul 2012 - Feb 2018 - Research Scientist

- Deep online/offline question-answering models. I led the first deep sequence scoring model launches in Google's Web QA system.
- Large scale precomputation of factoid and non-factoid question-answer pairs from query logs and Web documents.
- Enabling flexibly structured (html) answers in Google's answer box.
- KG schema and semantic parser induction from logs, KG, and Web QA system.
- Model-based knowledge graph confidence estimation, and error detection.
- Knowledge graph construction from Web documents.

Carnegie Mellon University, Pittsburg

Jul 2006 - Jun 2012 - Graduate Research Assistant

- Efficient randomized reasoning approach for IR, NLP and recommendation
- Relational CRFs structure learning, and hidden variable induction
- Architecture the CMU's cross-lingual question answering system (JAVELIN) which was a precursor to the IBM Watson system

Microsoft Research Asia, Beijing

Jul 2003 - Jun 2006 - Graduate Research Assistant

- Large scale clustering and classification of online products using product image, text and other metadata
- Personalized search ranking by simulating user experience from search log
- An innovative learning to rank method that fit piecewise linear curves
- Automatic operating system troubleshooting based on text descriptions, system config changes, and system call events

State Key Lab of Intelligent Technology and Systems, Beijing

Feb 2001 - Jun 2003 - Undergraduate Research Assistant

- I was one of the main developers for the Tsinghua Aeolus soccer system, which was the world champion of RoboCup Simulation League in 2001 & 2002 using methods such as dynamic programming and geometry computation.
- I also design the debugger visualizing the players' world models, and the bandwidth constrained communication scheme between players.

EDUCATION

Carnegie Mellon University, Pittsburg

July 2006 - June 2012 - PhD in Language Technology

Thesis: Efficient Random Walk Inference with Knowledge Bases.

Tsinghua University, Beijing

July 2003 - June 2006 - Master in Computer Science

Thesis: Data Mining Problems in Automatic Computer Diagnosis.

July 1999 - June 2003 - Bachelor in Electronic Engineering

Thesis: Mining Spatial-Temporal Data Using Constructive Induction. (RoboCup)

SELECTED PUBLICATIONS

- Zhangyu Wang*, et al, Ni Lao*, Gengchen Mai*, **LocDiffusion: Identifying Locations on Earth by Diffusing in the Hilbert Space**. (in submission)
- Gengchen Mai*, Zeping Liu*, Ni Lao*, et al, Stefano Ermon, **SSIF: Physics-Inspired Implicit Representations for Spatial-Spectral Image Super-Resolution**. (in submission)
- Z Wang*, G Mai*, K Janowicz, N Lao*, **MC-GTA: Metric-Constrained Model-Based Clustering using Goodness-of-fit Tests with Autocorrelations**. [ICML 2024](#)
- Gengchen Mai*, et al, Ni Lao*, **Sphere2Vec: A general-purpose location representation learning over a spherical surface for large-scale geospatial predictions**. [ISPRS 2023](#)
- Gengchen Mai*, Ni Lao*, et al, Stefano Ermon, **CSP: Self-Supervised Contrastive Spatial Pre-Training for Geospatial-Visual Representations**, [ICML 2023](#)
- Luyu Gao, et al, Ni Lao*, Hongrae Lee*, Da-Cheng Juan*, Kelvin Guu*, **Rarr: Researching and revising what language models say, using language models**, [ACL 2023](#)
- Gengchen Mai*, et al, Ni Lao*, **Multi-Scale Representation Learning for Spatial Feature Distributions using Grid Cells**. [ICLR 2020](#)
- Fan Yang, et al, Ni Lao*, **Learning to Organize Knowledge with N-Gram Machines**. [ICLR 2018](#)
- Chen Liang*, et al, Ni Lao*, **Memory augmented policy optimization for program synthesis and semantic parsing**. [NeurIPS 2018](#)
- Chen Liang*, et al, Ni Lao*, **Neural symbolic machines: Learning semantic parsers on freebase with weak supervision**. [ACL 2017](#)
- Ni Lao, Tom Mitchell, William Cohen, **Random walk inference and learning in a large scale knowledge base**, [EMNLP 2011](#)
- N Lao, WW Cohen, **Relational retrieval using a combination of path-constrained random walks**, [Machine learning 2010](#)
- Yao Jinyi, Lao Ni, etc, Sun Zengqi, Technical solutions of TsinghuAeolus for robotic soccer. [RoboCup 2003](#)