

Ni Lao, PhD

HEADLINE

A world class expert in machine learning (ML), knowledge graph (KG), and natural language understanding (NLU), who wants to redefine what is possible in products and research. A strong technical manager who can lead teams in the design and integration of cutting edge research and create real business value.

EXPERIENCE

Mosaix.ai, Palo Alto

Feb 2018 - PRESENT - Co-founder and Chief Scientist

- Founded, built and ran a team of highly skilled NLP/ML engineers/researchers from the ground up.
- Defines the company's AI strategy and roadmap.
- Manages the development and scaling of NLP, ML and quality infrastructures as the foundation for the Mosaix Language AI platform.
- Engages with the research community and publishes at top AI conferences each year with cutting edge innovation.

Google, Mountain View

Jul 2012 - Feb 2018 - Senior Research Scientist

- Research and productization of NLU, KG and ML technologies.
(Tech lead for various projects in production)
- Deep question-answer scoring models for both online and offline data.
- Precomputing factoid and non-factoid question answer pairs from anticipated questions and web documents.
- Enabling flexibly structured (html) answers in Google's answer box.
- Semantic parser grammar induction from logs, KG, and search results.
- Model-based knowledge graph confidence estimation, and error detection
- Knowledge graph construction from Web documents.

Carnegie Mellon University, Pittsburgh

Jul 2006 - Jun 2012 - Research Assistant

- Random walk inference on knowledge graphs for 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.
- Chinese Semantic Role Labeling. Recurrent Net for constituency parsing
- Automated evaluation for utility-based information distillation

Microsoft Research Asia, Beijing

Jul 2003 - Jun 2006 - Research Assistant

- Large scale clustering and classification of online products using product image, text and other metadata.
- Analyzed personalized search strategies by simulating user experience from search logs.
- Experiment with learning to rank methods for Bing search performance.
- Automatic OS system troubleshooting based on text description, system config change, and system call events.

State Key Lab of Intelligent Technology and Systems, Beijing

Feb 2001 - Jun 2003 - Research Assistant

- Was the main developer of the Tsinghua Aeolus robo-soccer system.
- Our system was the world champion of the RoboCup competition Simulation League in 2001 and 2002.

EDUCATION

Carnegie Mellon University, Pittsburg

July 2006 - June 2012 - PhD in Language Technology

Thesis: Efficient Random Walk Inference with Knowledge Bases.

Advisor: William W. Cohen, Teruko Mitamura

Tsinghua University, Beijing

July 2003 - June 2006 - Master in Computer Science

Thesis: Data Mining Problems in Automatic Computer Diagnosis.

Advisor: Ji-Rong Wen, Chunping Li

Tsinghua University, Beijing

July 1999 - June 2003 - Bachelor in Electronic Engineering

Thesis: Mining Spatial-Temporal Data Using Constructive Induction.

Advisor: Zeng-Qi Sun, Chunping Li

PRESENTATIONS

2019, A Review of Google's Lingvo. A review requested by Synced

2019, Yann LeCun's Cake Explained. based on my interview with Robin.ly

2019, Weakly Supervised Natural Language Understanding. JiangMen online tutorial

2018, Weakly Supervised Natural Language Understanding. AIFrontiers tutorial

2018, Do Androids Dream of Great Success? Medium

2017, Neural Symbolic Language Understanding. Stanford Deep Learning Seminar

2017, Text Generation Survey. Google Language AI, Boston

2017, \w Xipeng Qiu, Knowledge Acquisition. CCKS lectures, Chengdu, China

2016, NIPS 2016 Overview. Google Search Quality

2015, New Development in Knowledge Acquisition, Inference, and Applications. Lectures at China Computer Federation (CCF), Beijing, China

2012, Elephant and AI. LTI Colloquium

2012, Programming by Demonstrations and Verbal Commands. LTI Colloquium

2011, Beyond Shallow Semantics. LTI Colloquium

2011, CCG, Fractal, and Emergence. LTI Colloquium

2011, Reinforcement Learning In An Unknown Domain, CMU

2010, Probabilistic Ontology Model. LTI Colloquium

2009, Split-Emit Process for Natural Language Generation. Adv. NLP seminar CMU

2007, Read The Web. Advanced IR seminar, CMU

2007, Schema Extraction Model. Advanced IR seminar, CMU

2007, Knowledge Acquisition From Text--A Survey, CMU

ACADEMIC SERVICES

I am one of the recipients of AI for Earth Microsoft Azure Compute Grant in 2020 for my proposal on Deep Species Spatio-Temporal Distribution Modeling for Biodiversity Hotspot Prediction.

I co-organized the Deep RL Meets Structured Prediction workshop at ICLR 2019.

(Reviews & Chairs)

2020: ACL, AACL, CoNLL, EMNLP, IJCAI, SIGIR

2019: ACL, AACL, CCL, CoNLL, EMNLP, ICLR, IJCAI, NAACL, SIGIR, TKDE

2018: ACL, CCKS, COLING, EMNLP, NAACL, NLPCC, NIPS, SIGIR

2017: ACL, CCKS, EMNLP, IJCAI, IJCNLP, SIGIR, TKDE, WSDM, Google Research Grants

2016: CIKM, COLING, IJCAI, NAACL, TKDE, WWW, Google Research Grants

2015: CIKM, ICML, IJCAI, MLJ, NIPS, TKDE

2014: CIKM, KDD

2013: CIKM, Google Research Grants

2012: JCST, KDD, Neurocomputing J., SIGIR, TKDD, PKDD

2011: NIPS, T. Fuzzy Systems

2010: CIKM

SELECTED PUBLICATIONS

50+ Papers, 5 Patents, 3300+ citations

(Patents) [Google Patents](#)

Cheng HE, Ni Lao, Xiuqi Tan, Sumang Liu, Method and apparatus for searching historical data, US20190370398A1, 2019

Ni Lao, Chen Liang, Quoc V Le, John Blitzer, Neural question answering system, US20190130251A1. 2019

Ni Lao, Lukasz Mieczyslaw Kaiser, Nitin Gupta, Afroz Mohiuddin, Preyas Popat, Answer to question neural networks, US20180114108A1. 2018

Ni Lao, Jiazhong NIE, Fan Yang, Natural language processing with an n-gram machine, W02019083519A1, 2017

Kevyn B Collins-thompson, Ni Lao, Context-Aware Query Alteration, US20120233140A1, 2012

(Peer reviewed papers) [Google Scholar](#)

Gengchen Mai, Krzysztof Janowicz, Bo Yan, Rui Zhu, Ling Cai, Ni Lao, Multi-Scale Representation Learning for Spatial Feature Distributions using Grid Cells, in ICLR 2020. Spotlight Presentation.

Mai et al, Ni Lao Semantically-Enriched Search Engine for Geoportals: A Case Study with ArcGIS Online, In: Proceedings of AGILE 2020, Jun. 16 - 19, 2020, Chania, Crete, Greece.

Mai et al, Ni Lao Contextual Graph Attention for Answering Logical Queries over Incomplete Knowledge Graphs, In the 10th ACM International Conference on Knowledge Capture (K-CAP 2019)

Felix Wu, Boyi Li, Lequn Wang, Ni Lao, John Blitzer, Kilian Q. Weinberger, Integrated Triaging for Fast Reading Comprehension, Preprint, 2019.

Jacob Biloki, Chen Liang, Ni Lao, Neural Program Planner for Structured Predictions, In ICLR 2019, Workshop on Deep Reinforcement Learning Meets Structured Prediction.

Felix Wu, Boyi Li, Lequn Wang, Ni Lao, John Blitzer, Kilian Q. Weinberger, FastFusionNet: New State-of-the-Art for DAWNbench SQuAD, Technical Report, 2019

Gengchen Mai, Krzysztof Janowicz and Cheng He, Sumang Liu, Ni Lao. POIReviewQA: A Semantically Enriched POI Retrieval and Question Answering Dataset, In 12th Workshop on Geographic Information Retrieval (GIR 2018)

Chen Liang, Mohammad Norouzi, Jonathan Berant, Quoc Le, Ni Lao, Memory Augmented Policy Optimization for Program Synthesis with Generalization, In NIPS 2018

T. Mitchell, et al, Never-Ending Learning, Communications of the ACM, 2018

Juanzi Li, Ming Zhou, Guilin Qi, Ni Lao, Tong Ruan, Jianfeng Du, Knowledge Graph and Semantic Computing. Language, Knowledge, and Intelligence, Communications in Computer and Information Science, Springer, 2017

Fan Yang, Jiazhong Nie, William W. Cohen, Ni Lao, Learning to Organize Knowledge with N-Gram Machines, ICLR 2018 Workshop. AKBC best poster award.

Chen Liang, Jonathan Berant, Quoc Le, Kenneth D. Forbus, Ni Lao, Neural Symbolic Machines: Learning Semantic Parsers on Freebase with Weak Supervision, ACL 2017.

Felix Wu, Ni Lao, John Blitzer, Guandao Yang, and Kilian Weinberger, Fast Reading Comprehension with Convnets, Technical Report, 2017.

Ni Lao, Einat Minkov and William Cohen, Learning relational features with backward random walks, ACL 2015.

William Yang Wang, Kathryn Mazaitis, Ni Lao, Tom M. Mitchell, William W. Cohen, Efficient Inference and Learning in a Large Knowledge Base: Reasoning with Extracted Information using a Locally Groundable First-Order Probabilistic Logic, Machine Learning Journal (MLJ 2015), Springer.

T. Mitchell, et al, (2015): Never-Ending Learning in AAAI-2015.

Ni Lao, Jun Zhu: Contrastive Feature Induction for Efficient Structure Learning of Conditional Random Fields. CoRR abs/1406.7445 (2014) code

Dong et al, Knowledge Vault: A Web-Scale Approach to Probabilistic Knowledge Fusion. KDD, 2014

Ni Lao, Amarnag Subramanya, Fernando Pereira, William W. Cohen Reading The Web with Learned Syntactic-Semantic Inference Rules. EMNLP, 2012

Ni Lao, William W. Cohen, Personalized Reading Recommendations for Saccharomyces Genome Database. DILS, 2012

Ni Lao, Tom Mitchell, William W. Cohen, Random Walk Inference and Learning in A Large Scale Knowledge Base. EMNLP, 2011

Jun Zhu, Ni Lao, Ning Chen, Eric P. Xing Conditional Topical Coding: an Efficient Topic Model Conditioned on Rich Features. KDD, 2011

Ni Lao, William W. Cohen, Relational retrieval using a combination of path-constrained random walks Machine Learning, 2010, Volume 81, Number 1, Pages 53-67 ECML, 2010

Ni Lao, Jun Zhu, Liu Liu, Yandong Liu, William W. Cohen, Efficient Relational Learning with Hidden Variable Detection. NIPS, 2010

Ni Lao, William W. Cohen, Fast Query Execution for Retrieval Models based on Path Constrained Random Walks. KDD, 2010

Jun Zhu, Ni Lao, E. P. Xing, Grafting-Light: Fast, Incremental Feature Selection and Structure Learning of Markov Random Fields. KDD, 2010

Lao, Ni, Hideki Shima, Teruko Mitamura and Eric Nyberg. 2008. Query Expansion and Machine Translation for Robust Cross-Lingual Information Retrieval, in Proceedings of NTCIR-7 Workshop, Japan.

Shima, Hideki, Ni Lao, Eric Nyberg and Teruko Mitamura. 2008. Complex Cross-lingual Question Answering as Sequential Classification and Multi-Document Summarization Task, in Proceedings of NTCIR-7 Workshop, Japan.

W. Zuo, N. Lao, Y. Geng, and K. Ma. 2008. GeoSVM: an efficient and effective tool to predict species' potential distributions. Journal of Plant Ecology, 1(2): 143-145.

Yiming Yang, Abhimanyu Lad, Ni Lao, Abhay Harpale, Bryan Kisiel, Monica Rogati, Utility-based information distillation over temporally sequenced documents, SIGIR, pp. 31-38, 2007.

Chun Yuan; Ni Lao; Ji-Rong Wen; Jiwei Li; Zheng Zhang; Yi-Min Wang; Wei-Ying Ma, Automated Known Problem Diagnosis with Event Traces, EuroSys, 2006.

Ni Lao, Ji-Rong Wen, Wei-Ying Ma, Yi-Min Wang, Combine High Level Symptom and Low Level State Information for Configuration Fault Diagnosis, LISA, 2004.

Ji-Rong Wen, Ni Lao, Wei-Ying Ma, Probabilistic Model for Contextual Retrieval, SIGIR, 2004.

Archana Ganapathi, Yi-Min Wang, Ni Lao, Ji-Rong Wen, Why PCs Are Fragile and What We Can Do About It: A Study of Windows Registry Problems, Dependable System and Network (DSN), 2004.

Jinyi Yao, Lao Ni, Fan Yang, Yunpeng Cai, Zengqi Sun, Technical Solutions of TsinghuAeolus for Robotic Soccer. Robocup 2003: 205-213, RoboCup, pp. 205-213, 2003