

Xin Su, Ph.D.

CONTACT INFORMATION	Phone: (503) 383-4379 Email: xinsu626@gmail.com	
EDUCATION	University of Arizona, Tucson, AZ Ph.D. in Information, Advisor: Steven Bethard Loyola University Chicago, Chicago, IL M.S. in Computer Science M.S. in Supply Chain Management Western Oregon University, Monmouth, OR B.S. in Business	Dec. 2024 May 2020 Jun. 2016
INDUSTRY & RESEARCH EXPERIENCE	Intel Labs, USA <i>AI Research Scientist</i> - Submitted 7 papers within 5 months of joining (3 to NeurIPS, 2 accepted at CVPR 2025 workshops, 1 to ACL ARR, 1 accepted at ACL 2025 workshop) covering multimodal reasoning, LLM adaptation, temporal reasoning, retrieval-augmented generation, and synthetic data generation - Leading a 5-person team in designing multimodal agents for complex computer use tasks, driving cross-functional collaboration with multiple business units for Intel AI PC deployment - Architecting deep research-style structured multimodal agents for business intelligence, enabling complex reasoning over databases to generate comprehensive analytical reports - Achieved exceptional performance recognition within first quarter of joining Intel Labs, USA <i>AI Research Intern</i> - Developed SK-VQA synthetic dataset with 2M+ examples using novel generation methods to enhance vision-language models' external knowledge integration and multimodal RAG performance; resulted in first-author ICML 2025 Spotlight Oral paper (top 1%). - Received Tech Excellence Award for outstanding research contributions and customer engagement. - Proposed novel KG generation method and RAG approach, trained specialized LLMs, scaled inference to 512 Gaudi 2 cards on Intel Tiber AI Cloud using Kubernetes; served as Intel's primary contributor in strategic customer collaboration, leading to co-first author ACL ARR submission. - Proposed novel chain-of-thought method enabling foundation models to effectively leverage diverse knowledge sources for complex reasoning in knowledge-intensive tasks; demonstrated significant improvements across both open-source	Jan. 2025 - Present May - Aug. 2022, May 2023 - Dec. 2024

and proprietary foundation models; resulted in **first-author NAACL 2024 publication**.

- Proposed novel temporal graph construction and transformer fusion method that significantly improved temporal reasoning capabilities in LLMs; resulted in **first-author EMNLP 2023 Findings publication**.

Computational Language Understanding Lab, University of Arizona 2020-2024
Graduate Research Associate

- Proposed novel compositional annotation schema-based time normalization method to anchor temporal expressions to timelines, resulting in an open-source Python library; led to first-author NeurIPS 2025 submission.
- Built transformer-based constrained generation methods for semantic parsing, automatically translating natural language to structured query languages (SQL and Cypher).
- Systematically analyzed and compared source-free domain adaptation methods using self-training, active learning, and data augmentation without access to source domain data; resulted in **first-author ACL 2022 publication**.

Loyola Natural Language Processing Lab, Loyola University Chicago 2018-2020
Graduate Research Assistant

- Designed hierarchical transformer framework for encoding long clinical documents using pre-trained language models; implementation integrated into open-source Python library (cnlp_transformers).
- Pre-trained BERT model on 10+ years of electronic health records from Loyola Medical Center for clinical domain adaptation.
- Collaborated with medical professionals to develop ML/NLP solutions for clinical applications, including patient cohort identification, semantic similarity, and medical concept extraction from 10+ years of electronic health records.

PUBLICATIONS

Peer-Reviewed Publications

- [1] **Xin Su**, Sungduk Yu, Phillip Howard, Steven Bethard. A Semantic Parsing Framework for End-to-End Time Normalization. *The Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS)*. 2025
- [2] **Xin Su***, Man Luo*, Kris W Pan, Tien Pei Chou, Vasudev Lal, Phillip Howard. SK-VQA: Synthetic Knowledge Generation at Scale for Training Context-Augmented Multimodal LLMs (**Spotlight Oral, top 1%**). *International Conference on Machine Learning (ICML)*. 2025
- [3] **Xin Su**, Phillip Howard, Steven Bethard. Transformer-Based Temporal Information Extraction and Application: A Review. *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2025
- [4] **Xin Su**, Man Luo, David Cobbley, Shachar Rosenman, Vasudev Lal, Phillip Howard,

Vision-based Assistants in the Real-World. Plan-Action-Reflection: A Three-Role Agentic Framework For Computer Use Agent Task. *Vision-based Assistants in the Real-World at The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2025

- [5] Sezen Perçin, **Xin Su**, Qutub Sha Syed, Phillip Howard, Aleksei Kuvshinov, Leo Schwinn, Kay-Ulrich Scholl, Investigating the Robustness of Retrieval-Augmented Generation at the Query Level. *Generation, Evaluation and Metrics (GEM) at the Association for Computational Linguistics (ACL)*. 2025
- [6] Man Luo, David Cobbley, **Xin Su**, Shachar Rosenman, Vasudev Lal, Shao-Yen Tseng, Phillip Howard, DPO Learning with LLMs-Judge Signal for Computer Use Agents. *The 4th Workshop on Computer Vision in the Wild (CVinW) at The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2025
- [7] Neale Ratzlaff, Man Luo, **Xin Su**, Vasudev Lal, Phillip Howard, Training-Free Mitigation of Language Reasoning Degradation After Multimodal Instruction Tuning. *Proceedings of the AAAI Symposium Series*. 2024
- [8] **Xin Su**, Tiep Le, Steven Bethard, Phillip Howard, Semi-Structured Chain-of-Thought: Integrating Multiple Sources of Knowledge for Improved Language Model Reasoning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL): Human Language Technologies 2024*
- [9] **Xin Su**, Phillip Howard, Nagib Hakim, Steven Bethard. Fusing Temporal Graphs into Transformers for Time-Sensitive Question Answering. *Findings of the Association for Computational Linguistics: EMNLP 2023*
- [10] **Xin Su**, Yiyun Zhao, Steven Bethard. A Comparison of Strategies for Source-Free Domain Adaptation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*. 2022.
- [11] **Xin Su**, Yiyun Zhao, Steven Bethard. The university of arizona at semeval-2021 task 10: Applying self-training, active learning and data augmentation to source-free domain adaptation. In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval)*. 2021.
- [12] Egoitz Laparra, **Xin Su**, Yiyun Zhao, Ozlem Uzuner, Timothy Miller, Steven Bethard. Semeval-2021 task 10: Source-free domain adaptation for semantic processing. In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval)*. 2021.
- [13] Sujay Kulshrestha, Dmitriy Dligach, **Xin Su**, Richard Gonzalez, Cara Joyce, Matthew M Churpek, Majid Afshar. Classification of Chest Injury Severity Using Clinical Documents. *American Medical Informatics Association (AMIA) Informatics Summit*. 2021. (peer-reviewed abstract)

- [14] Anoop Mayampurath, Matthew Churpek, **Xin Su**, Sameep Shah, Elizabeth Munroe, Bhakti Patel, Dmitriy Dligach, Majid Afshar. External Validation of an Acute Respiratory Distress Syndrome Prediction Model Using Radiology Reports. *Critical Care Medicine*. 2020.
- [15] **Xin Su**, Timothy Miller, Majid Afshar, Dmitriy Dligach. Learning Hierarchical Transformer-based Representations of Clinical Notes. *American Medical Informatics Association (AMIA) Symposium*. Chicago, IL. November 2020. (peer-reviewed abstract)
- [16] **Xin Su**, Anoop Mayampurath, Matthew Churpek, Sameep Shah, Bhakti Patel, Dmitriy Dligach, Majid Afshar. External Validation of an Acute Respiratory Distress Syndrome Prediction Model Using Clinical Text. *American Thoracic Society (ATS) International Conference 2020*. Philadelphia, Pennsylvania, May 2020.
- [17] **Xin Su**, Timothy Miller, Farig Sadeque, Majid Afshar, Dmitriy Dligach. Using Transformer-based Approaches for Measuring Semantic Similarity. *National NLP Clinical Challenges (N2C2) Workshop at AMIA 2019 Annual Symposium*. Washington, D.C., November 2019. (peer-reviewed abstract)

Preprints and Under Review

- [1] **Xin Su**, Kadir Bulut Ozler, Egoitz Laparra, Alon Geva, Yingya Li, Timothy A. Miller, Guergana K Savova, Phillip Howard, Steven Bethard, How many examples does it take for fine-tuning to outperform few-shot prompting? A study of medical tasks and domain adaptation. *Under Review*. 2025
- [2] Man Luo, Sungduk Yu, **Xin Su**, Vasudev Lal, Majid Afshar, Phillip Howard, A Diverse Autoregressive Benchmark for Next-Day Progress Note Generation and Evaluation. *Under Review*. 2025
- [3] Prafulla Kumar Choubey*, **Xin Su***, Man Luo*, Xiangyu Peng, Caiming Xiong, Tiej Le, Shachar Rosenman, Vasudev Lal, Phil Mui, Ricky Ho, Phillip Howard, Chien-Sheng Wu. Distill-SynthKG: Distilling Knowledge Graph Synthesis Workflow for Improved Coverage and Efficiency. *Under Review*. 2025
- [4] **Xin Su**, Timothy Miller, Xiyu Ding, Majid Afshar, Dmitriy Dligach. Classifying Long Clinical Documents with Pre-trained Transformers. *arXiv:2105.06752*. 2021.

* denotes equal contribution.

TEACHING

ISTA 322: Data Engineering. University of Arizona, Teaching Associate, Spring 2023.
 ISTA 457 / INFO 557 Neural Networks. University of Arizona, Teaching Associate, Fall 2022.

SERVICE

Association for Computational Linguistics Rolling Review (ARR). Reviewer. 2023-present.

Forty-second International Conference on Machine Learning (ICML). Reviewer. 2025.

The Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS). Reviewer. 2025.

The Thirteenth International Conference on Learning Representations (ICLR). Reviewer. 2024

The 39th Annual AAAI Conference on Artificial Intelligence. Program Committee (AAAI). 2024.

Forty-first International Conference on Machine Learning (ICML). Reviewer. 2024.

The Twelfth International Conference on Learning Representations (ICLR). Reviewer. 2024.

Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS). Reviewer. 2023.

The Conference on Empirical Methods in Natural Language Processing (EMNLP). Program Committee. 2023.

The 61st Annual Meeting of the Association for Computational Linguistics (ACL). Program Committee. 2023.

The Conference on Empirical Methods in Natural Language Processing (EMNLP). Program Committee. 2022.

American Medical Informatics Association (AMIA) Symposium. Reviewer. 2022.

American Medical Informatics Association (AMIA) Symposium. Reviewer. 2020.

TECHNICAL SKILLS

Languages: Python, Java, C, C++, R, SQL, Bash, JavaScript, $\text{\LaTeX}$

Tools & Software: Pytorch, Tensorflow, Keras, Scikit-learn, cTAKES, Git, Docker, Microsoft Azure, Kubernetes, Slurm, Linux

HONORS & AWARDS

College of Information Science Dean's List of Distinguished Scholars, University of Arizona 2024

Edsger W. Dijkstra High Achievement Award (CS), Loyola University Chicago 2020

Travel Award to attend N2C2 Workshop at AMIA, Loyola University Chicago 2019

Beta Gamma Sigma, Loyola University Chicago 2018

Merit Scholarship, Loyola University Chicago 2016

Dean's Honor Roll, Western Oregon University 2016