

Revanth Gangi Reddy

✉ g.revanthreddy111@gmail.com • 📄 [gangiswag.github.io](https://github.com/gangiswag)

Interests: Large Language Models, Retrieval and Reasoning, Natural Language Processing
Citizenship: Indian

Education

PhD in Computer Science	University of Illinois at Urbana-Champaign	GPA - 4.0/4.0	Aug'22 - May'25
MS in Computer Science	University of Illinois at Urbana-Champaign	GPA - 3.85/4.0	Jan'21 - July'22
B. Tech, Computer Science	Indian Institute of Technology (IIT), Madras	GPA - 9.16/10.0	July'14 - May'18

Internships

Salesforce AI Research **Spring 2025**
Research Intern, RAG team *Palo Alto, United States*

- Working on training agents for complex information-seeking Retrieval-Augmented Generation tasks.

Apple **Summer 2023**
Research Intern, Knowledge Platform team *Seattle, United States*

- Worked on complex question answering over knowledge graphs for Siri's Knowledge Platform.

Allen Institute for Artificial Intelligence **Summer 2022**
Research Intern, AllenNLP team *Seattle, United States*

- Worked on pre-training strategies for domain adaptation of neural information retrieval models.

Amazon Alexa **Spring 2022**
Applied Scientist Intern, Alexa Web Information team *Seattle, United States*

- Worked on summarizing reported speech about events in news and introduced a new benchmark, SumREN, with an automatic training data generation strategy.

IBM Research AI **Summer 2018**
Research Intern, Watson Conversations team *New Delhi, India*

- Worked on better neural memory architectures for improving the performance of task-oriented dialog systems.

Work Experience

Alexa SocialBot Grand Challenge 5 **Oct 2022 - Aug 2023**
Team Lead, Advisor: Prof. ChengXiang Zhai *Champaign, United States*

- Building a **social chatbot** that can converse coherently and engagingly with humans on popular topics and current events. Awarded a grant of **\$250,000** to lead a team of 13 students from UIUC.

University of Illinois at Urbana-Champaign **Oct 2020 - Present**
Research Assistant, Advisor: Prof. Heng Ji *Champaign, United States*

- Working on projects in the areas of multimodal question answering, knowledge retrieval and claim detection in news articles. Ongoing collaborations with Amazon, IBM Research, Columbia University and UNC Chapel Hill.

IBM Research AI **Oct 2019 - Oct 2020**
AI Resident, Advisors: Dr. Avi Sil, Dr. Salim Roukos *New York, United States*

- Worked on projects in the areas of question answering, open-domain knowledge retrieval and AMR parsing. Contributed to five research papers (four first-author) and three patents during the AI Residency program.

Microsoft **Oct 2018 - Sep 2019**
Software Engineer, Data Integration team *Vancouver, Canada*

- Worked in the team that develops data connectors for [PowerApps](#), [LogicApps](#) and [Flow](#). Built the entire data pipeline for a scalable Azure-based multi-region logging infrastructure that currently handles 7 billion logs a day.

Recent Research Projects

Inference-time Re-ranker Relevance Feedback for Neural IR [PDF](#) **Under Review**

Revanth Reddy, Pradeep Dasigi, Arafat Sultan, Arman Cohan, ... , Hannaneh Hajishirzi

- Proposed a novel **inference-time re-ranker feedback** mechanism that improves the recall of retrieval in multidomain, multilingual and multimodal settings in a retrieve-and-rerank framework.

SmartBook: AI-Assisted Situation Report Generation [PDF](#) **Under Review**

Revanth Reddy, Yi Fung, Qi Zeng, ..., Paul Sullivan, Heng Ji

- Introduced SmartBook, a generalizable automated framework designed to **assist human analysts** in structured real-time situation report generation from large news corpora.

Research Publications

CoRNStack: High-Quality Contrastive Data for Better Code Ranking [PDF](#) **ICLR 2025**

Tarun Suresh*, Revanth Reddy*, Yifei Xu, ... , Brandon Duderstadt, Heng Ji

- Released a large-scale, high-quality contrastive training dataset for code that leads to **state-of-the-art performance across a variety of code retrieval tasks**.

INFOGENT: An Agent-Based Framework for Web Information Aggregation [PDF](#) **NAACL 2025 Findings**

Revanth Reddy*, Sagnik Mukherjee*, Jeonghwan Kim*, Zhenhailong Wang*, ... , Heng Ji

- Introduced a novel modular and **feedback-driven framework for web information aggregation** involving three distinct components: Navigator, Extractor and Aggregator.

AGRaME: Any-Granularity Ranking with Multi-Vector Embeddings [PDF](#) **EMNLP 2024**

Revanth Reddy*, Omar Attia*, Yunyao Li, Heng Ji, Saloni Potdar

- Introduced the concept of Any-Granularity Ranking which leverages multi-vector embeddings to **rank at varying levels of granularity** while maintaining encoding at a single (coarser) level of granularity.

FIRST: Faster Improved Listwise Reranking with Single Token Decoding [PDF](#) **EMNLP 2024**

Revanth Reddy*, JaeHyeok Doo*, Yifei Xu*, ... , Avirup Sil, Heng Ji

- Proposed a **novel listwise LLM reranking approach leveraging the output logits** of the first generated identifier to directly obtain a ranked ordering of the candidates.

Progressive Responses for ChatBots with Real-Time Internet Search [PDF](#) **WSDM 2024 Demo**

Revanth Reddy, Sharath Chandra, Hao Bai, ..., ChengXiang Zhai

- Introduced the use of **progressive response generation to integrate real-time web search results** into voice-based chatbots, thereby reducing user wait times by upto 50%.

Commonsense-Guided Search Query Generation for Open-Domain Dialog [PDF](#) **Findings of EMNLP 2023**

Revanth Reddy, Hao Bai, Wentao Yao, ..., Heng Ji, ChengXiang Zhai

- To tackle passive conversations, we propose to integrate **social commonsense reasoning** to guide the generation of search queries in knowledge-powered conversations.

SumREN: Summarizing Reported Speech about Events in News [PDF](#) **AAAI 2023**

Revanth Reddy, Heba Elfardy, Hou Pong Chan, Kevin Small, Heng Ji

- Proposed a new challenging task of summarizing reported speech about events in news with a **multi-document summarization benchmark**, SUMREN, that leverages GPT-3 for automatic silver-training data generation.

NewsClaims: Claim Detection from News with Attribute Knowledge [PDF](#) **EMNLP 2022**

Revanth Reddy, Sai Chetan, Zhenhailong Wang, Yi Fung, ... , Kevin Small, Heng Ji

- Proposed a new benchmark for **knowledge-aware claim detection**, that re-defines the claim detection problem to include extraction of additional background attributes related to the claim.

A Zero-Shot Claim Detection Framework using Question Answering [PDF](#) **COLING 2022**

Revanth Reddy, Sai Chinthakindi, Yi R. Fung, Kevin Small, Heng Ji

- Proposed a **fine-grained claim detection framework** that leverages zero-shot question answering using directed questions to solve a diverse set of sub-tasks such as topic filtering, claim object detection and claimer detection.

Entity-Conditioned Question Generation for Robust Attention Distribution in Neural IR [PDF](#) *SIGIR 2022*

Revanth Reddy, Arafat Sultan, Martin Franz, Avi Sil, Heng Ji

- Proposed an **entity-conditioned data augmentation** strategy that generates questions about sparsely-attended entities in the passage, to help improve neural IR models by learning to distribute attentions over the passage.

Towards Robust Neural Retrieval Models with Source Domain Synthetic Pre-Training [PDF](#) *COLING 2022*

Revanth Reddy, Vikas Yadav, Arafat Sultan, Martin Franz, Vittorio Castelli, Heng Ji, Avi Sil

- Improved the **zero-shot performance** of state-of-the-art neural IR models on both **in-domain** and **out-of-domain** datasets by pre-training with **synthetic questions** generated automatically from raw text passages.

MuMuQA: Multimedia Multi-Hop News Question Answering via Cross-Media Grounding [PDF](#) *AAAI 2022*

Revanth Reddy, Xilin Rui, Manling Li, ... , Mohit Bansal, Avi Sil, Shih-Fu Chang, Heng Ji

- Proposed a new benchmark for **multimodal** question answering over news articles, with a novel data generation framework for generating questions that are **grounded on objects in images**.

COVID-19 Claim Radar: A Structured Claim Extraction and Tracking System [PDF](#) *ACL Demo 2022*

Manling Li, Revanth Reddy, Ziqi Wang, Yi-Shyuan Chiang, ... , Zixuan Zhang, Heng Ji

- Built a system that **automatically extracts claims** in news articles and shows a **comprehensive structured view** of such claims, with rich attributes (such as claimers and their affiliations) and associated knowledge elements.

Synthetic Target Domain Supervision for Open Retrieval QA [PDF](#) *SIGIR 2021*

Revanth Reddy, Bhavani Iyer, Arafat Sultan, ... , Vittorio Castelli, Radu Florian, Salim Roukos

- Leveraged an **automatic text-to-text generation** idea to improve the performance of state-of-the-art open-domain **end-to-end** question answering systems in a specialized domain, such as COVID-19.

Multi-Stage Pre-training for Low-Resource Domain Adaptation [PDF](#) *EMNLP 2020*

Revanth Reddy*, Rong Zhang*, Arafat Sultan, ... , Avi Sil, Todd Ward, Radu Florian, Salim Roukos

- Proposed **synthetic pre-training** objectives by using structure in unlabeled text, that can transfer to downstream tasks with considerable gains in the **IT Domain**.

Answer Span Correction in Machine Reading Comprehension [PDF](#) *Findings of EMNLP 2020*

Revanth Reddy, Arafat Sultan, Rong Zhang, Efsun Kayi, Vittorio Castelli, Avi Sil

- Proposed an approach for **correcting partial match answers** ($EM=0, 0<F1<1$) into exact match ($EM=1, F1=1$) and obtained upto **1.3%** improvement in both monolingual and multilingual evaluation.

Pushing the Limits of AMR Parsing with Self-Learning [PDF](#) *Findings of EMNLP 2020*

Revanth Reddy*, Young-suk Lee*, Ramon Astudillo*, Tahira Naseem*, ... , Salim Roukos

- Proposed **self-learning approaches** via generation of synthetic text and synthetic AMR as well as refinement of actions from the oracle, achieving **state-of-the-art** performance on benchmark AMR 1.0 and AMR 2.0 datasets.

Patents

Method for Answer Span Correction [PDF](#)

Improving Model Performance through Text-to-Text Transformation via Distant Supervision [PDF](#)

Scholastic Achievements

- Awarded the Siebel Scholarship for the class of 2022. ([Link](#))
- All India Rank - 127 in JEE Advanced 2014, taken by more than 1.3 million students.
- Shortlisted among the top 32 in India from engineering stream for Aditya Birla Scholarships 2014.
- All India Rank - 44 in KVPY 2012, taken by close to 200,000 students.
- Topped the Chennai center in Mimamsa 2017 prelims, an All India Science Quiz held by IISER Pune.

Extra-Curricular Activities

- Won bronze medal in badminton in Schroeter (Inter-hostel Sports Tournament) 2016 at IIT Madras.
- Reviewed papers as part of the program committee for multiple top-tier academic conferences.
- Taught mathematics to middle-school students in Chennai as a part of National Social Service's Math Teach 2016.