

Panupong (Ice) Pasupat

ppasupat@cs.stanford.edu | (+1) 857-919-5187 | <https://ppasupat.github.io>

Education

Stanford University

Ph.D. in Computer Science

Advisor: Percy Liang

Dissertation Title: Natural Language Interfaces for Semi-Structured Web Pages

Stanford, CA

2013–2019

Massachusetts Institute of Technology

B.S. in Electrical Engineering and Computer Science

B.S. in Mathematics (double major)

Cambridge, MA

2009–2013

Work Experience

Google DeepMind

Research Scientist

Mountain View, CA

2019–present

- Conducted research on semantic parsing, retrieval-augmented models, few-shot learning, and factuality.
- Led a collaboration with Google Assistant to develop a query understanding model that can patch urgent errors without retraining.
- Led a research team to develop an automated evidence-based fact checker, now powering Google Gemini’s “Double-check Response” and “Sources and Related Content” features.
- Supervised Ph.D. research interns.

Facebook Conversational AI, Facebook

Research Intern

Menlo Park, CA

2018

- Improved a neural shift-reduce model for parsing sentences into a hierarchical semantic representation.
- Analyzed errors and designed novel top-down and bottom-up parsing algorithms to address the errors.

Google Research, Google

Software Engineering Intern

Mountain View, CA

2015

- Developed deep learning models for paraphrase detection.
- Proposed negative sampling methods using linguistic resources to better distinguish closely related words.

Speech and Dialog Research Group, Microsoft Research

Research Intern

Mountain View, CA

2014

- Bootstrapped classifiers for detecting knowledge base relations in spoken queries in an unsupervised fashion.
- Mined queries from search engine query click logs and automatically labeled relations using distant supervision from knowledge graphs.

Natural Language Processing Lab, Tokyo Institute of Technology

Exchange Student

Yokohama, Japan

2013

- Experimented on Tweet sentiment analysis using different classifiers and features.
- Applied structural correspondence learning to incorporate unlabeled data.

Spoken Language Systems Group, MIT CSAIL Lab

Cambridge, MA

Researcher Intern

2012–2013

- Designed web interfaces on Amazon Mechanical Turk to collect spoken sentences and their semantic labeling.
- Trained sequence tagging models by implementing features for conditional random fields, resulting in English and Chinese models for categorizing words in speech queries.
- Deployed the models in speech-enabled mobile applications for movie, flight, and restaurant recommendations.

Dropbox Inc.

San Francisco, CA

Engineering Intern

2011

- Wrote scripts to periodically analyze the usage pattern of Dropbox users. Investigated methods to optimize data calculation and caching for the analytics team.
- Implemented bug filtering and email notification system in the error log viewer, making critical errors get noticed and fixed faster.

Reflective Commonsense Thinking, MIT Media Lab

Cambridge, MA

Researcher Intern

2010

- Implemented algorithms to solve error-correcting subgraph isomorphism and analogical matching problems.
- Applied the algorithms on analogy problems, enabling the system to make decisions based on past experiences.

Language of Thought, MIT Department of Linguistics

Cambridge, MA

Researcher Intern

2010

- Designed algorithms to automatically measure vowel formant frequencies from sound files in order to observe the patterns and constraints of vowels in spoken languages.
- Designed experiments on Amazon Mechanical Turk to study sound constraints during language acquisition.

Publications

Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities

arXiv preprint, 2025

LOFT: Scalable and More Realistic Long-Context Evaluation

(former title: Can Long-Context Language Models Subsume Retrieval, RAG, SQL, and More?)

Jinhyuk Lee, Anthony Chen, Zhuyun Dai, Dheeru Dua, Devendra Singh Sachan, Michael Boratko, Yi Luan, Séb Arnold, Vincent Perot, Siddharth Dalmia, Hexiang Hu, Xudong Lin, Panupong Pasupat, Aida Amini, Jeremy R Cole, Sebastian Riedel, Iftekhar Naim, Ming-Wei Chang, Kelvin Guu

North American Chapter of the Association for Computational Linguistics (NAACL) Findings, 2025

In-context learning with retrieved demonstrations for language models: A survey

Man Luo, Xin Xu, Yue Liu, Panupong Pasupat, Mehran Kazemi

Transactions on Machine Learning Research (TMLR), 2024

Large language models as analogical reasoners

Michihiro Yasunaga, Xinyun Chen, Yujia Li, Panupong Pasupat, Jure Leskovec, Percy Liang, Ed H. Chi, Denny Zhou

International Conference on Learning Representations (ICLR), 2024

Retrieval-Augmented Parsing for Complex Graphs by Exploiting Structure and Uncertainty

Zi Lin, Quan Yuan, Panupong Pasupat, Jeremiah Z. Liu, Jingbo Shang

Empirical Methods in Natural Language Processing (EMNLP) Findings, 2023

From Pixels to UI Actions: Learning to Follow Instructions via Graphical User Interfaces

Peter Shaw, Mandar Joshi, James Cohan, Jonathan Berant, Panupong Pasupat, Hexiang Hu, Urvashi Khandelwal, Kenton Lee, Kristina Toutanova
Conference on Neural Information Processing Systems (NeurIPS), 2023

PURR: Efficiently Editing Language Model Hallucinations by Denoising Language Model Corruptions

Anthony Chen, Panupong Pasupat, Sameer Singh, Hongrae Lee, Kelvin Guu
arXiv preprint, 2023

Dr. ICL: Demonstration-Retrieved In-context Learning

Man Luo, Xin Xu, Zhuyun Dai, Panupong Pasupat, Mehran Kazemi, Chitta Baral, Vaiva Imbrasaite, Vincent Y Zhao
arXiv preprint, 2023

On Compositional Uncertainty Quantification for Seq2seq Graph Parsing

Zi Lin, Du Phan, Panupong Pasupat, Jeremiah Zhe Liu, Jingbo Shang
International Conference on Learning Representations (ICLR), 2023

RARR: Researching and Revising What Language Models Say, Using Language Models

Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Y. Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, Kelvin Guu
Association for Computational Linguistics (ACL), 2023

Evaluating the Impact of Model Scale for Compositional Generalization in Semantic Parsing

Linlu Qiu, Peter Shaw, Panupong Pasupat, Tianze Shi, Jonathan Herzig, Emily Pitler, Fei Sha, Kristina Toutanova
Empirical Methods in Natural Language Processing (EMNLP), 2022

Meta-Learning Fast Weight Language Models

Kevin Clark, Kelvin Guu, Ming-Wei Chang, Panupong Pasupat, Geoffrey Hinton, Mohammad Norouzi
Empirical Methods in Natural Language Processing (EMNLP), 2022

Generate-and-Retrieve: Use Your Predictions to Improve Retrieval for Semantic Parsing

Yury Zemlyanskiy, Michiel de Jong, Joshua Ainslie, Panupong Pasupat, Peter Shaw, Linlu Qiu, Sumit Sanghai, Fei Sha
International Conference on Computational Linguistics (COLING), 2022

Improving Compositional Generalization with Latent Structure and Data Augmentation

Linlu Qiu*, Peter Shaw*, Panupong Pasupat, Pawel Nowak, Tal Linzen, Fei Sha, Kristina Toutanova
North American Chapter of the Association for Computational Linguistics (NAACL), 2022

Controllable Semantic Parsing via Retrieval Augmentation

Panupong Pasupat, Yuan Zhang, Kelvin Guu
Empirical Methods in Natural Language Processing (EMNLP), 2021

Graph-Based Decoding for Task Oriented Semantic Parsing

Jeremy Cole, Nanjiang Jiang, Panupong Pasupat, Luheng He, Peter Shaw
Empirical Methods in Natural Language Processing (EMNLP) Findings, 2021

QA-Driven Zero-shot Slot Filling with Weak Supervision Pretraining

Xinya Du, Luheng He, Qi Li, Dian Yu, Panupong Pasupat, Yuan Zhang
Association for Computational Linguistics (ACL), 2021

Unlocking Compositional Generalization in Pre-trained Models Using Intermediate Representations

Jonathan Herzig, Peter Shaw, Ming-Wei Chang, Kelvin Guu, Panupong Pasupat, Yuan Zhang
arXiv preprint, 2021

Few-shot Intent Classification and Slot Filling with Retrieved Examples

Dian Yu, Luheng He, Yuan Zhang, Xinya Du, Panupong Pasupat, Qi Li
North American Chapter of the Association for Computational Linguistics (NAACL), 2021

Compositional Generalization and Natural Language Variation: Can a Semantic Parsing Approach Handle Both?

Peter Shaw, Ming-Wei Chang, Panupong Pasupat, Kristina Toutanova
Association for Computational Linguistics (ACL), 2020

REALM: Retrieval-Augmented Language Model Pre-Training

Kelvin Guu*, Kenton Lee*, Zora Tung, Panupong Pasupat, Ming-Wei Chang
International Conference on Machine Learning (ICML), 2020

SPoC: Search-based Pseudocode to Code

Sumith Kulal*, Panupong Pasupat*, Kartik Chandra, Mina Lee, Oded Padon, Alex Aiken, Percy Liang
Conference on Neural Information Processing Systems (NeurIPS), 2019

Span-based Hierarchical Semantic Parsing for Task-Oriented Dialog

Panupong Pasupat, Sonal Gupta, Karishma Mandyam, Rushin Shah, Mike Lewis, Luke Zettlemoyer
Empirical Methods on Natural Language Processing (EMNLP), 2019

Improving Semantic Parsing for Task Oriented Dialog

Arash Einolghozati, Panupong Pasupat, Sonal Gupta, Rushin Shah, Mrinal Mohit, Mike Lewis, Luke Zettlemoyer
Conversational AI Workshop at NeurIPS, 2018

Mapping Natural Language Commands to Web Elements

Panupong Pasupat, Tian-Shun Jiang, Evan Liu, Kelvin Guu, Percy Liang
Empirical Methods on Natural Language Processing (EMNLP), 2018

Reinforcement Learning on Web Interfaces using Workflow-Guided Exploration

Evan Zheran Liu*, Kelvin Guu*, Panupong Pasupat*, Tianlin Shi, Percy Liang
International Conference on Learning Representations (ICLR), 2018

Macro Grammars and Holistic Triggering for Efficient Semantic Parsing

Yuchen Zhang, Panupong Pasupat, Percy Liang
Empirical Methods on Natural Language Processing (EMNLP), 2017

From Language to Programs: Bridging Reinforcement Learning and Maximum Marginal Likelihood

Kelvin Guu, Panupong Pasupat, Evan Zheran Liu, Percy Liang
Association for Computational Linguistics (ACL), 2017

Inferring Logical Forms From Denotations

Panupong Pasupat, Percy Liang
Association for Computational Linguistics (ACL), 2016

Simpler Context-Dependent Logical Forms via Model Projections

Reginald Long, Panupong Pasupat, Percy Liang
Association for Computational Linguistics (ACL), 2016

Compositional Semantic Parsing on Semi-Structured Tables

Panupong Pasupat, Percy Liang
Association for Computational Linguistics (ACL), 2015

Unsupervised Relation Detection Using Automatic Alignment of Query Patterns Extracted from Knowledge Graphs and Query Click Logs

Panupong Pasupat, Dilek Hakkani-Tür
Interspeech, 2015

Zero-Shot Entity Extraction from Web Pages

Panupong Pasupat, Percy Liang
Association for Computational Linguistics (ACL), 2014

Query Understanding Enhanced by Hierarchical Parsing Structures

Jingjing Liu, Panupong Pasupat, Yining Wang, Scott Cyphers, Jim Glass
Automatic Speech Recognition and Understanding Workshop (ASRU), 2013

ASGARD: a Portable Architecture for Multilingual Dialogue Systems

Jingjing Liu, Panupong Pasupat, Scott Cyphers, Jim Glass
International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2013

A Conversational Movie Search System Based on Conditional Random Fields

Jingjing Liu, Scott Cyphers, Panupong Pasupat, Ian Mcgraw, Jim Glass
Interspeech, 2012

Patents

Machine-Learned User Interface Command Generator Using Pretrained Image Processing Model

Peter Thomas Shaw, Mandar Joshi, Kristina Nikolova Toutanova, James Fischl Cohan, Jonathan Haim Berant, Kenton Chiu Tsun Lee, Panupong Pasupat, Hexiang Hu, Urvashi Khandelwal
US Patent, 2025

Retrieval-Augmented Language Model Pre-Training and Fine-Tuning

Kenton Chiu Tsun Lee, Kelvin Gu, Zora Tung, Panupong Pasupat, Ming-Wei Chang
US Patent, 2021

Semantic Entity Relation Detection Classifier Training

Dilek Hakkani-Tür, Panupong Pasupat
US Patent, 2019

Honors and Awards

Putnam Mathematical Competition

Ranked 23rd 2009

International Olympiad in Informatics (IOI)

Gold Medal (Ranked 2nd) 2008

International Mathematical Olympiad (IMO)

Gold Medal 2007
2 Silver Medals 2005, 2006

Teaching Experience

Information Retrieval and Web Search (CS 276)

Head Course Assistant **Stanford University**
Spring 2016

Artificial Intelligence: Principles and Techniques (CS 221)

Head Course Assistant **Stanford University**
Fall 2014

Services

Conference Area Chair

- Association for Computational Linguistics (ACL) 2023
- European Chapter of the Association for Computational Linguistics (EACL) 2023
- International Conference on Computational Linguistics (COLING) 2022
- Empirical Methods on Natural Language Processing (EMNLP) 2021

Journal Reviewer

- Transactions of the Association for Computational Linguistics (TACL) 2022–present

Conference Reviewer (Program Committee Member)

- Association for the Advancement of Artificial Intelligence (AAAI) 2020
- Association for Computational Linguistics (ACL) 2016, 2017, 2018, 2020, 2021
- Conference on Automated Knowledge Base Construction (AKBC) 2019, 2020, 2021
- International Conference on Computational Linguistics (COLING) 2020
- Conference on Language Modeling (COLM) 2025
- Conference on Natural Language Learning (CoNLL) 2019
- European Chapter of the Association for Computational Linguistics (EACL) 2021
- Empirical Methods on Natural Language Processing (EMNLP) 2015, 2017, 2019, 2020, 2022
- International Conference on Learning Representations (ICLR) 2020
- North American Chapter of the Association for Computational Linguistics (NAACL) 2019, 2021
- Joint Conference on Lexical and Computational Semantics (*SEM) 2017, 2018, 2019

Rolling Review Reviewer

- ACL Rolling Review (2 month cycles for top-tier conferences: ACL, EACL, NAACL, and EMNLP) 2024, 2025

Workshop Reviewer

- Machine Reading for Question Answering (MRQA) 2021
- Natural Language Processing for Programming (NLP4Prog) 2021
- Table Representation Learning (TRL) 2024, 2025
- Widening Natural Language Processing (WiNLP) 2020

Website Chair

- North American Chapter of the Association for Computational Linguistics (NAACL) 2021, 2022

Coordinator and Grader

- International Mathematical Olympiad (IMO) 2015
- USA Mathematical Olympiad (USAMO) 2011, 2013

Skills

Programming: fluent in Python (TensorFlow, PyTorch) and JavaScript; working knowledge of Java and C++

Languages: Thai (native), English (fluent), Japanese (intermediate), Chinese (beginner)