

NATHANIEL WEIR

nweir127.github.io · nweir127@gmail.com

EDUCATION

- 2019–2024 **Johns Hopkins University**
Ph.D. in Computer Science
Advisor: Benjamin Van Durme
Thesis: *Compositional Neuro-Symbolic Reasoning over Natural Language*
- 2019–2021 **Johns Hopkins University**
M.S.E. in Computer Science
Advisor: Benjamin Van Durme
- 2015–2019 **Brown University**
Sc.B. in Applied Mathematics & Computer Science, *magna cum laude* with Honors (GPA: 3.95/4)
Advisors: Ugur Cetintemel, Carsten Binnig, and Ellie Pavlick
Thesis: *Bootstrapping Generalization in Neural Text-to-SQL Semantic Parsing Models*

RESEARCH INTERESTS

Neuro-Symbolic Reasoning, Logical Verification, Natural Language Inference, Large Language Models, Knowledge Grounding, Explainable AI, Information Retrieval, Question Answering

RESEARCH EXPERIENCE

- Fall 2024 **Amazon Web Services**
– Present *Applied Scientist II*
- Fall 2019 **Center for Language and Speech Processing at Johns Hopkins University**
– Summer 2024 *PhD Researcher*
Research includes retrieval-augmented, language model-based neuro-symbolic reasoning, constrained language generation, and improving robustness of question answering models.
- Summer 2023 **Allen Institute for Artificial Intelligence**
– Winter 2023 *PhD Research Intern – Aristo Reasoning Team*
Mentor: Peter Clark
Researched extracting and reasoning over generalizable scientific microtheories from language models. Also led projects on code-based instruction tuning and abductive situational reasoning.
- Summer 2022 **Microsoft Semantic Machines**
PhD Research Intern
Mentor: Harsh Jhamtani
Worked on knowledge-constrained dialogue generation. Published a benchmark dataset of real video game dialogue trees annotated with granular ontology constraints.
- Summer 2021 **Microsoft Research – Montreal**
PhD Research Intern – Deep Learning and Language Team
Mentors: Harm Van Seijen, Xingdi Yuan and Marc-Alexandre Côté
Explored hierarchical language-based planning in embodied agents. Published a benchmark test-
ing for learning compositional tasks in a Minecraft-like domain.
- Spring 2017 **Brown University**
– May 2019 *Undergraduate Researcher*
Researched domain adaptation methods for text-to-SQL neural semantic parsing.

PREPRINTS

- 2026 Vikash Singh, Darion Cassel, **Nathaniel Weir**, Nick Feng, Sam Bayless. [VERGE: Formal Refinement and Guidance Engine for Verifiable LLM Reasoning](#).
- 2026 Jiefu Ou, Sapana Chaudhary, Kaj Bostrom, **Nathaniel Weir**, Shuai Zhang, Huzefa Rangwala, George Karypis. [MaxCode: A Max-Reward Reinforcement Learning Framework for Automated Code Optimization](#).
- 2025 Sam Bayless, Stefano Buliani, Darion Cassel, Byron Cook, ..., **Nathaniel Weir**, Michael W. Whalen, Jianan Yao. [A Neurosymbolic Approach to Natural Language Formalization and Verification](#).

PUBLICATIONS

- 2026 Yu Feng, **Nathaniel Weir**, Kaj Bostrom, Sam Bayless, Darion Cassel, Sapana Chaudhary, Benjamin Kiesl-Reiter, Huzefa Rangwala. [VeriCoT: Neuro-symbolic Chain-of-Thought Validation via Logical Consistency Checks](#). ICLR 2026.
- 2025 Zhengping Jiang, Jingyu Zhang, **Nathaniel Weir**, Seth Ebner, Miriam Wanner, Kate Sanders, Daniel Khashabi, Anqi Liu, Benjamin Van Durme. [Core: Robust Factual Precision Scoring with Informative Sub-Claim Identification](#). Findings of ACL 2025.
- 2025 **Nathaniel Weir**, Bhavana Dalvi Mishra, Orion Weller, Oyvind Tafjord, Sam Hornstein, Alexander Sabol, Peter Jansen, Benjamin Van Durme, and Peter Clark. [From Models to Microtheories: Distilling a Model’s Topical Knowledge for Grounded Question Answering](#). ICLR 2025.
- 2025 Dongwei Jiang, Jingyu Zhang, Orion Weller, **Nathaniel Weir**, and Benjamin Van Durme, Daniel Khashabi. [SELF-\[IN\]CORRECT: LLMs Struggle with Refining Self-Generated Responses](#). AAAI 2025.
- 2024 **Nathaniel Weir***, Muhammad Khalifa*, Linlu Qiu, Orion Weller, and Peter Clark. [Learning to Reason via Program Generation, Emulation, and Search](#). NeurIPS 2024.
- 2024 **Nathaniel Weir**, Kate Sanders, Orion Weller, Shreya Sharma, Dongwei Jiang, Zhengping Zhang, Bhavana Dalvi Mishra, Oyvind Tafjord, Peter Jansen, Peter Clark, and Benjamin Van Durme. [Enhancing Systematic Decompositional Natural Language Inference Using Informal Logic](#). EMNLP 2024.
- 2024 Kate Sanders, **Nathaniel Weir**, and Benjamin Van Durme. [TV-TREES: Multimodal Entailment Trees for Neuro-Symbolic Video Reasoning](#). EMNLP 2024.
- 2024 **Nathaniel Weir**, Ryan Thomas, Randolph d’Amore, Kellie Hill, Benjamin Van Durme, and Harsh Jhamtani. [Ontologically Faithful Generation of Non-Player Character Dialogues](#). EMNLP 2024.
- 2024 **Nathaniel Weir**, Peter Clark, and Benjamin Van Durme. [NELLIE: A Neuro-Symbolic Inference Engine for Grounded, Compositional, and Explainable Reasoning](#). IJCAI 2024.
- 2024 Xinrui Zou, Ming Zhang, **Nathaniel Weir**, Benjamin Van Durme, and Nils Holzenburger. [Reframing Tax Law Entailment as Analogical Reasoning](#). Jurix Special Workshop on AI, Law and Philosophy.
- 2024 Orion Weller, Marc Marone, **Nathaniel Weir**, Dawn Lawrie, Daniel Khashabi, and Benjamin Van Durme. [“According to ...” Prompting Language Models Improves Quoting from Pre-Training Data](#). EACL 2024.
- 2024 Orion Weller, Aleem Khan, **Nathaniel Weir**, Dawn Lawrie, and Benjamin Van Durme. [Defending Against Poisoning Attacks in Open-Domain Question Answering](#). EACL 2024.
- 2023 **Nathaniel Weir**, Xingdi Yuan, Marc-Alexandre Côté, Matthew J. Hausknecht, Romain Laroche, Ida Momennejad, Harm van Seijen and Benjamin Van Durme. [One-Shot Learning from a Demonstration with Hierarchical Latent Language](#). AAMAS 2023.

- 2021 Jiefu Ou*, **Nathaniel Weir***, Anton Belyy*, Felix Yu, and Benjamin Van Durme. [InFillmore: Frame-Guided Language Generation with Bidirectional Context](#). StarSem 2023.
- 2020 **Nathaniel Weir**, João Sedoc, and Benjamin Van Durme. [COD3S: Diverse Generation with Discrete Semantic Signatures](#). EMNLP 2020.
- 2020 **Nathaniel Weir**, Adam Poliak, and Benjamin Van Durme. [Probing Neural Language Models for Human Tacit Assumptions](#). CogSci 2020 **Oral Presentation**.
- 2020 **Nathaniel Weir**, Prasetya Utama, Alex Galakatos, Andrew Crotty, Amir Ilkhechi, Shekar Ramaswamy, Rohin Bhusan, Nadja Geisler, Benjamin Hattasch, Steffen Eger, Ugur Cetintemel, and Carsten Binnig. [DBPal: A Fully Pluggable NL2SQL Training Pipeline](#). SIGMOD 2020.
- 2018 Fuat Basik, Benjamin Hattasch, Amir Ilkhechi, Arif Usta, Shekar Ramaswamy, Prasetya Utama, **Nathaniel Weir**, Carsten Binnig and Ugur Cetintemel. [DBPal: A Learned NL-Interface for Databases](#). SIGMOD 2019 Demo.
- 2017 Prasetya Utama, **Nathaniel Weir**, Carsten Binnig, and Ugur Cetintemel. [Voice-based Data Exploration: Chatting with your Database](#). SCAI 2017.

AWARDS

- 2021 - 2024 **NSF Graduate Research Fellowship**
- 2019 **CRA Outstanding Undergraduate Researcher Award, Honorable Mention**
- 2019 **SIGMOD Undergraduate Research Competition, 2nd place**
- 2019 **Brown Computer Science Senior Award**
- 2019 **Sigma Xi Electee**
- 2019 **Brown CS Undergraduate Research Symposium, 3rd Place**
- 2018 **Randy F. Pausch CS Undergraduate Summer Research Award, \$10,000**
- 2017 **Karen T. Romer Undergraduate Teaching and Research Award, \$3,500**

PRESENTATIONS

Neuro-Symbolic Entailment Tree Search in the Era of Retrieval-Augmented LLMs

- 07/2024 Talk Microsoft Research
- 04/2024 Talk Amazon Web Services
- 03/2024 Talk Kensho

Ontologically Faithful Generation of Non-Player Dialogues

- 07/2023 Talk DialDoc Workshop at ACL 2023

Language Models as Proposal Functions in a Neuro-Symbolic Expert System

- 05/2023 Talk Massachusetts Institute of Technology
- 05/2023 Talk Brown University

One-Shot Learning from a Demonstration with Hierarchical Latent Language

- 06/2023 Poster AAMAS 2023
- 03/2022 Talk 9th Mid-Atlantic Student Colloquium on Speech, Language and Learning

Probing Neural Language Models for Human Tacit Assumptions

- 03/2020 Poster 8th Mid-Atlantic Student Colloquium on Speech, Language and Learning

DBPal: A Fully Pluggable Natural Language Interface to Databases

01/2019	Talk	North East Database Day @ MIT
10/2018	Talk	IBM AI Systems Day @ MIT
01/2018	Demo	North East Database Day @ MIT

TECHNICAL SKILLS

Languages Python (expert), R, Bash, Prolog, C/C++, SQL, Java, MATLAB, Scala

Tools PyTorch, HuggingFace Transformers, LangChain, fairseq, faiss, vLLM, DeepSpeed, MySQL, Jupyter, RStudio, NLTK, Mechanical Turk

TEACHING EXPERIENCE

Teaching Assistant, Johns Hopkins University

Fall 2022 CS 601.470/670: Artificial Agents. Instructor: Benjamin Van Durme (co-taught)

Teaching Assistant, Brown University

Fall 2018 CSCI1570: Design and Analysis of Algorithms. Instructor: Paul Valiant

Spring 2018 CSCI0220: Discrete Structures and Probability. Instructor: Caroline Klivans

Spring 2017 CSCI0220: Discrete Structures and Probability. Instructor: Caroline Klivans

SERVICE

Reviewing:

ACL Rolling Review (2022, 2023, 2024), ICLR 2025, NeurIPS 2024, EMNLP 2023, KnowledgeNLP-AAAI'23 ICLR 2021 (Secondary), ACL 2020 (Secondary), AKBC 2020 (Secondary)

Committees:

JHU CS Ph.D. Admissions Committee (2021-2024)

JHU CLSP Application Support Program for underrepresented students (2022-2023)

MENTORING

- Shreya Sharma (JHU MS, 2023-2024)
- Dongwei Jiang (JHU MS, joint with Orion Weller, 2023-2024)
- Sandipan Majhi (JHU MS, 2023)
- Jingyu Zhang (JHU BS, 2022-2023)
- Xiao Ye (JHU MS, 2022)
- Chenyu Zhang (JHU BS, 2020-2021)
- Jiefu Ou (HKUST BS, 2020-2021)
- Wei Liu (UIUC BS, 2020)