

# Eylon Caplan

Email: [ecaplan@purdue.edu](mailto:ecaplan@purdue.edu)

Website: [eyloncaplan.github.io](https://eyloncaplan.github.io)

GitHub: [@eyloncaplan](https://github.com/eyloncaplan)

LinkedIn: [@eyloncaplan](https://www.linkedin.com/in/eyloncaplan)

EDUCATION	<p><b>Purdue University</b>, West Lafayette, Indiana, USA <span style="float: right;"><b>(expected) 2027</b></span> Ph.D. in Natural Language Processing, Department of Computer Science <span style="float: right;"><b>3.8/4.0</b></span></p> <p><b>University of Nebraska–Lincoln</b>, Lincoln, Nebraska, USA <span style="float: right;">2023</span> B.S. in Computer Science and Mathematics (Minors: Physics, Spanish) <span style="float: right;"><b>3.99/4.0</b></span></p>
SKILLS	<p>AI, Machine Learning, Natural Language Processing, LLMs, VLMs, Information Retrieval, RAG, Reinforcement Learning, Big Data, Multiprocessing, Benchmarking, Software Engineering, Deep Learning, Generative AI</p> <p>Python, Jupyter, Pandas, NumPy, PyTorch, HuggingFace, Transformers, LangChain, ColBERT, pyserini, BERTopic, Dask, Docker, FAISS, Flask, Git, Hydra, vLLM, Kubernetes, Imdeploy, SLURM, Scripting</p>
PUBLICATIONS	<p><b>CONCEPTCARVE: Dynamic Realization of Evidence</b> Eylon Caplan and Dan Goldwasser ACL 2025 Main Conference Poster, Vienna, Austria on July 26–August 2</p> <p><b>VIBE: Can a VLM Read the Room?</b> Tania Chakraborty, Eylon Caplan, and Dan Goldwasser Findings of EMNLP 2025 Poster, Suzhou, China on November 5–9</p> <p><b>SPLITS! Flexible Sociocultural Linguistic Investigation at Scale</b> Eylon Caplan, Tania Chakraborty, and Dan Goldwasser ACL 2026 Main Conference Oral, San Diego, California on July 2–7</p> <p><b>TAIGR: Towards Modeling Influencer Content on Social Media via Structured, Pragmatic Inference</b> Nishanth Sridhar Nakshatri, Eylon Caplan, Rajkumar Pujari, and Dan Goldwasser ACL 2026 Main Conference Poster, San Diego, California on July 2–7</p>
RESEARCH EXPERIENCE	<p><b>Graduate Research Assistant, Purdue NLP Lab</b> (Advisor: Dan Goldwasser) <span style="float: right;">2023–Present</span> <i>Developing NLP frameworks to model and improve social reasoning in large-scale online communities.</i></p> <ul style="list-style-type: none"><li>- Built <b>CONCEPTCARVE</b>, a framework uniting <b>LLM reasoning</b> with <b>scalable retrieval</b>, reranking, and clustering to capture abstract concepts manifesting in social communities, achieving a 26.03% relative improvement over LLM keyword expansion.</li><li>- Engineered <b>SPLITS!</b>, a <b>9.7M-post dataset</b> and sociolinguistic sandbox to systematically evaluate LLM-generated hypotheses. Developed an <b>automated pipeline</b> utilizing <i>statistical lift</i> and <i>BERT embeddings</i> to surface non-trivial phenomena, reducing manual review by 15–18x.</li><li>- Exposed the “Visual Social-Pragmatic Inference gap” where <b>VLMs</b> misinterpret social visual cues, such as a sad smile. To measure this, created <b>VIBE</b>, a 994-instance video benchmark isolating this <i>reasoning failure</i>, and explored <b>frontier model performance</b>, demonstrating humans outperform VLMs by &gt;17%.</li></ul>
INDUSTRY EXPERIENCE	<p><b>Software Engineering Intern, Hudl</b> <span style="float: right;">2022– 2023</span> Developed and deployed a CV pipeline using PyTorch to perform OCR on basketball scoreboards from live video. Integrated the service into production environment, with real-time overlays for live streaming on HudlTV.</p>
TEACHING AND CURRICULUM DESIGN	<p><b>Course Developer, Purdue University</b> <span style="float: right;">2023– 2025</span></p> <ul style="list-style-type: none"><li>- Designed a module and four-part project about the RAG pipeline for the <i>AI Forge</i> course. Project included parts teaching model inference, prompting, in-context learning, retrieval, and retrieval augmentation. Also designed an evaluation pipeline of student code on computing cluster.</li><li>- Designed assignments and course content for a new course, <i>Data Structures and Algorithms for AI</i>. Created four course projects, covering topics like trees, stacks, queues, big data hashing, fuzzy word search, and graphs.</li></ul>
KEY COURSES	<p><b>Graduate Level:</b> Advanced Topics in Reasoning with LLMs, NLP, Deep Learning, Reasoning about Programs</p>
KEY COURSE PROJECTS	<p><b>LLM Feedback for Proofs</b> <span style="float: right;">May 2023–Dec 2023</span> Tested various methods of injecting feedback from an LLM in order to generate correct symbolic proofs in the Isabelle proof solver for competition math problems. Course project for <i>Adv. Topics in Reasoning with LLMs</i>.</p> <p><b>Math Expression Style Transfer</b> <span style="float: right;">Aug 2024–Dec 2024</span> Developed an LLM BFS algorithm for converting math expressions into various simplified/expanded forms using only examples, with guaranteed equivalence. Course project for <i>Reasoning about Programs</i>.</p>
AWARDS	<p>Corporate Partners Scholarship <span style="float: right;">2023– 2024</span> Purdue Science Excellence Scholarship <span style="float: right;">2023– 2024</span></p>