

ARTEMIS PANAGOPOULOU

@artemisp@seas.upenn.edu [@artemisp.github.io](https://github.com/artemisp) [@linkedin.com/in/apanagop](https://www.linkedin.com/in/apanagop)
[@scholar.google.com/apanagopoulou](https://scholar.google.com/citations?user=apana) [@github.com/artemisp](https://github.com/artemisp)

EDUCATION

University of Pennsylvania, Philadelphia, PA

Doctor of Philosophy, Computer and Information Science

Aug, 2021 - Present

Research Interests: Natural Language Processing, Computer Vision

Advisors: Chris Callison-Burch, Mark Yatskar

GPA: 3.97/4.

Master of Science in Engineering, Computer and Information Science

Jan, 2018 - Aug, 2020 Thesis: *“Metaphor and Entailment: Looking at Metaphors Through the Lense of Textual Entailment”*

Advisor: Mitch Marcus GPA: 3.77/4.

Dual Degree in Artificial Intelligence

Aug, 2015 - Aug, 2020

Bachelor of Applied Science (BAS), Computer and Cognitive Science.

Thesis: *“Best-First-Model-Merge: From Theory to Implementation and Application”*

Advisor: Mitch Marcus

Bachelors of Arts (BA) Honors, Cognitive Science

Thesis: *“Optical Flow Estimation from Event Based Cameras Using Deep Spiking Neural Networks”*

Advisor: Kostas Daniilidis

Bachelors of Arts (BA) Honors, Philosophy

Thesis: *“On the suitability of Generative Difference Making for addressing challenges in Artificial Intelligence and Robotics.”*

Advisor: Lisa Miracchi

Minor in Mathematics

GPA: 3.59/4.

RESEARCH EXPERIENCE

Student Researcher

May, 2025 - Present

Google, Mountain View CA

- Work on Vision-Language Models for the Augmented Reality team.

Research Scientist Intern

May, 2023 - August, 2024

Salesforce, Palo-Alto CA

- Conducted multimodal AI research under the supervision of Dr. Juan Carlos Niebles.
- Led three projects: two involving models that combine vision, language, audio, and 3D, resulting to an ECCV publication and an ARR (ACL) submission; and one on Visual Programming accepted to CVPR.
- Contributed to a CVPR publication on 3D-vision-text contrastive foundation modeling.

Research Assistant

May, 2019 - May, 2020

General Robotics, Automation, and Sensing (GRASP) Lab, University of Pennsylvania

- Worked on estimating optical flow from event based cameras (supervised and unsupervised) using Spiking Neural Networks. (Supervisor: Prof. Kostas Daniilidis)

Research Assistant

May, 2019 - Aug, 2019

Kod*Lab, University of Pennsylvania

- Developed a simulation for physically parameterized soft bellow-shaped robots with multiple degrees of freedom. (Supervisor: Prof. Daniel Koditschek)

AWARDS AND FUNDING

CTL Graduate Fellowship for Teaching Excellence	August, 2024
AWS Research Funding for Fair and Trustworthy AI	May, 2023
Alexa Taskbot Competition Finalist	February, 2022
President's Engagement Prize	May, 2020
Dean's List	Aug, 2017 - May, 2020
Penn Engineering Exceptional Service Award	March, 2019

PUBLICATIONS

- Panagopoulou, Artemis, Honglu Zhou, Silvio Savarese, Caiming Xiong, Chris Callison-Burch, Mark Yatskar, Juan Carlos Niebles. “Visual Unit Tests for More Robust Visual Programming”, Accepted to Conference in Computer Vision and Pattern Recognition (2025)
- Panagopoulou, Artemis, Le Xue, Ning Yu, Junnan Li, Dongxu Li, Shafiq Joty, Ran Xu, Silvio Savarese, Caiming Xiong, Juan Carlos Niebles. “X-InstructBLIP: A Framework for Aligning X-Modal Instruction Aware Representations to LLMs and Emergent Cross-modal Reasoning”, Accepted to European Conference on Computer Vision (2024)
- Panagopoulou, Artemis*, Coby Melkin*, and Chris Callison-Burch. “Evaluating Vision-Language Models on Bistable Images.” In Proceedings of The 13th edition of the Workshop on Cognitive Modeling and Computational Linguistics (CMCL 2024) **Best Paper Award**
- Xue, Le, Ning Yu, Shu Zhang, Artemis Panagopoulou, Junnan Li, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, Silvio Savarese. “ULIP-2: Towards Scalable Multimodal Pre-training for 3D Understanding” Accepted to Conference in Computer Vision and Pattern Recognition (2024)
- Chakrabarty, Tuhin, Arkadiy Saakyan, Olivia Winn, Artemis Panagopoulou, Yue Yang, Marianna Apidianaki, and Smaranda Muresan. “I Spy a Metaphor: Large Language Models and Diffusion Models Co-Create Visual Metaphors.” In The 61st Annual Meeting Of The Association For Computational Linguistics. 2023.
- Yang, Yue, Artemis Panagopoulou, Shenghao Zhou, Daniel Jin, Chris Callison-Burch, and Mark Yatskar. “Language in a Bottle: Language Model Guided Concept Bottlenecks for Interpretable Image Classification.” Accepted to Conference in Computer Vision and Pattern Recognition (2023)
- Yue Yang*, Artemis Panagopoulou*, Marianna Apidianaki, Mark Yatskar and Chris Callison-Burch. “Visualizing the Obvious: A Concreteness-based Ensemble Model for Noun Property Prediction.” Findings of EMNLP 2022.
- Panagopoulou, Artemis, Manni Arora Li Zhang Dimitri Cugini, Weiqiu You, Yue Yang Liyang Zhou, Yuxuan Wang Zhaoyi Hou, Alyssa Hwang, Lara Martin, Sherry Shi Chris Callison-Burch, and Mark Yatskar. “QuakerBot: A household dialog system powered by large language models”, Alexa Prize TaskBot Challenge Proceedings (2022)
- Yue Yang, Artemis Panagopoulou, Qing Lyu, Li Zhang, Mark Yatskar, Chris Callison-Burch (2021). “Visual Goal-Step Inference using wikiHow.” EMNLP 2021 (Oral).
- Yang, Yue, Joongwon Kim, Artemis Panagopoulou, Mark Yatskar, and Chris Callison-Burch. “Induce, edit, retrieve: Language grounded multimodal schema for instructional video retrieval.” arXiv preprint arXiv:2111.09276 (2021)
- Chaney, Kenneth, Artemis Panagopoulou, Chankyu Lee, Kaushik Roy, and Kostas Daniilidis. “Self-supervised optical flow with spiking neural networks and event based cameras.” In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 5892-5899. IEEE, 2021.

PATENTS

- Systems and methods for multi-modal language models US/18400477
- Provisional Patent Application: Systems and Methods for Visual Programming (63/681,721)
- Provisional Patent Application: Systems and Methods for Training and Evaluating Multimodal Neural Network Based Language Models (63/656,510)

SELECTED INVITED TALKS

Vision Language Models Workshop <i>Women in Data Science Conference (WiDS), Philadelphia PA</i>	<i>Feb, 2025</i>
Bridging vision and language: Advances and Challenges <i>Princeton University, Princeton NJ</i>	<i>Dec, 2024</i>
Advancing Multimodal AI: Integrating Modalities, Tackling Complex Challenges, and Enhancing Interpretability <i>University of Pennsylvania, Philadelphia PA</i>	<i>Sep, 2024</i>

ACADEMIC SERVICE

Reviewed for **CVPR**, 2025

Reviewed for Student Research Workshop at the Nations of the Americas Chapter of the Association for Computational Linguistics (**NACL SRW**), 2025

Reviewed for **ACL ARR** (February, August 2024, February 2025)

Reviewed for Student Research Workshop at the Annual Meeting of the Association for Computational Linguistics (**ACL SRW**), 2023, 2024

Reviewed for The 61st Annual Meeting of the Association for Computational Linguistics (**ACL**), 2023

Reviewed for Multimodal Agents Workshop **ECCV**, 2024

Reviewed for **COLING**, 2025

TEACHING EXPERIENCE

Teaching Assistant Course: <i>CIS 530: Natural Language Processing</i> Instructor: <i>Prof. Mark Yatskar</i>	<i>Aug, 2024 - Dec, 2024, and Aug, 2022 - Dec, 2022</i>
Instructor <i>Prison Teaching Initiative at Princeton University, Southwoods State Prison</i> Instructors: <i>Artemis Panagopoulou, Joe Abatte, Uthsav Chitra</i>	<i>Sept, 2022 - Dec, 2022</i>
Elementary School Instructor <i>Python Coding Curriculum at Kohelet Yeshiva School (4-5 grade)</i> Instructor: <i>Artemis Panagopoulou</i>	<i>Aug, 2021 - May, 2022</i>
Teaching Assistant Course: <i>CIS 700: Interactive Fiction and Text Generation</i> Instructor: <i>Prof. Chris Callison-Burch, Dr. Lara Martin</i>	<i>Jan, 2022 - May, 2022</i>
Teaching Assistant Course: <i>CIS 521: Introduction to Artificial Intelligence</i> Instructor: <i>Prof. Chris Callison-Burch</i>	<i>Aug, 2021 - Dec, 2021</i>
Head Teaching Assistant Course: <i>MCIT 592: Mathematical Foundations of Computer Science</i> Instructor: <i>Prof. Val Tannen</i>	<i>Aug, 2018 - May, 2019</i>

Teaching Assistant

Jan, 2018 - May, 2018

Course: *CIS 262: Automata, Computability, and Complexity*

Instructor: *Dr. Nima Roohi*

LEADERSHIP AND ACTIVITIES

Advancing Women in Engineering (AWE), Board Member

August 2023 - Present

Association of Alumnae (AofA), Board Member

May 2023 - Present

Alexa Taskbot Competition [Finalist]

Aug 2021 - May 2022

Mind, Intelligence, Research, and Analysis (MIRA) Group

May 2018 - Aug 2019

Women in Computer Science (WiCS)

Jan 2019 - May 2019