Artemis Panagopoulou

Senior Research Scientist, Salesforce Research

artemis.panagop@gmail.com | artemisp.github.io
Google Scholar | GitHub: artemisp | LinkedIn: apanagop

Researcher in **multimodal AI, synthetic data, and evaluation**. I develop methods for learning and reasoning across images, video, audio, text, and 3D, with a focus on scalable data generation, spatial understanding, reliable visual reasoning, and interpretable models.

Research Experience

Senior Research Scientist *Sep 2026–Present*
Salesforce Inc., Salesforce Research, Palo Alto, CA

- Research and develop multimodal models for image, video, and text understanding, collaborating with engineering teams on product integration.
- Develop methods for evaluating visual interaction models in simulated environments.

Student Researcher *May 2025–Nov 2025*
Google, Augmented Reality, Mountain View, CA

- Developed MapTrace, a synthetic data generation pipeline for spatial reasoning and route tracing, with a public release of two million annotated map-path examples.

AI Research Intern *May 2023–Aug 2024*
Salesforce Inc., Salesforce Research, Palo Alto, CA

- Developed multimodal alignment and cross-modal reasoning methods spanning images, audio, video, and 3D, including **X-InstructBLIP** and **Contra4**.
- Developed visual unit testing for reliable visual reasoning (**ViUniT**) and contributed to scalable multimodal pretraining for 3D understanding (**ULIP-2**).

Education

University of Pennsylvania Philadelphia, PA
Ph.D., Computer and Information Science Aug 2021–Aug 2026
Thesis: *Synthesizing Missing Data for Multimodal Models*
Advisor: Mark Yatskar GPA: 3.97/4.0

M.S.E., Computer and Information Science Jan 2018–Aug 2020
Thesis: *Metaphor and Entailment: Looking at Metaphors Through the Lens of Textual Entailment*
Advisor: Mitch Marcus GPA: 3.77/4.0

Dual Degree in Artificial Intelligence Aug 2015–Aug 2020
B.A.S., Computer and Cognitive Science; B.A. (Honors), Cognitive Science;
B.A. (Honors), Philosophy. Minor in Mathematics. GPA: 3.59/4.0

B.A.S. thesis: Best-First-Model-Merge: From Theory to Implementation and Application. Advisor: Mitch Marcus.
Cognitive Science thesis: Optical Flow Estimation from Event Based Cameras Using Deep Spiking Neural Networks. Advisor: Kostas Daniilidis.
Philosophy thesis: On the suitability of Generative Difference Making for addressing challenges in Artificial Intelligence and Robotics. Advisor: Lisa Miracchi.

Publications and Preprints

Bold identifies my name; ∗ denotes equal contribution where indicated.

1. **Interactive Visual Gym: Evaluating Visual Interaction Models with Simulated Users in Interactive Worlds**
   Chenhao Zheng, Zixian Ma, Kai Li, Kanchana Ranasinghe, Jieyu Zhang, **Artemis Panagopoulou**, Weikai Huang, Jianxu Shangguan, Yuanyuan Tan, Silvio Savarese, Juan Carlos Niebles, Ranjay Krishna, Ran Xu, Honglu Zhou. *arXiv:2611.08286*, 2026.
2. **MapTrace: Scalable Data Generation for Route Tracing on Maps**
   Artemis Panagopoulou, Aveek Purohit, Achin Kulshrestha, Soroosh Yazdani, Mohit Goyal. *arXiv:2512.19609*, 2025.
3. **Contra4: Evaluating Contrastive Cross-Modal Reasoning in Audio, Video, Image, and 3D**
   Artemis Panagopoulou, Le Xue, Honglu Zhou, Silvio Savarese, Ran Xu, Caiming Xiong, Chris Callison-Burch, Mark Yatskar, Juan Carlos Niebles. *EMNLP*, 2025.
4. **ViUniT: Visual Unit Tests for More Robust Visual Programming**
   Artemis Panagopoulou, Honglu Zhou, Silvio Savarese, Caiming Xiong, Chris Callison-Burch, Mark Yatskar, Juan Carlos Niebles. *CVPR*, 2025.
5. **X-InstructBLIP: A Framework for Aligning Image, 3D, Audio, Video to LLMs and its Emergent Cross-modal Reasoning**
   Artemis Panagopoulou, Le Xue, Ning Yu, Junnan Li, Dongxu Li, Shafiq Joty, Ran Xu, Silvio Savarese, Caiming Xiong, Juan Carlos Niebles. *ECCV*, 2024.
6. **Evaluating Vision-Language Models on Bistable Images**
   Artemis Panagopoulou∗, Coby Melkin∗, Chris Callison-Burch. *Workshop on Cognitive Modeling and Computational Linguistics (CMCL)*, 2024. **Best Paper Award.**
7. **ULIP-2: Towards Scalable Multimodal Pre-training for 3D Understanding**
   Le Xue, Ning Yu, Shu Zhang, **Artemis Panagopoulou**, Junnan Li, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, Silvio Savarese. *CVPR*, 2024.
8. **I Spy a Metaphor: Large Language Models and Diffusion Models Co-Create Visual Metaphors**
   Tuhin Chakrabarty, Arkadiy Saakyan, Olivia Winn, **Artemis Panagopoulou**, Yue Yang, Marianna Apidianaki, Smaranda Muresan. *Findings of ACL*, 2023.
9. **Language in a Bottle: Language Model Guided Concept Bottlenecks for Interpretable Image Classification**
   Yue Yang, **Artemis Panagopoulou**, Shenghao Zhou, Daniel Jin, Chris Callison-Burch, Mark Yatskar. *CVPR*, 2023.
10. **Visualizing the Obvious: A Concreteness-based Ensemble Model for Noun Property Prediction**
    Yue Yang∗, **Artemis Panagopoulou**∗, Marianna Apidianaki, Mark Yatskar, Chris Callison-Burch. *Findings of EMNLP*, 2022.
11. **QuakerBot: A Household Dialog System Powered by Large Language Models**
    Artemis Panagopoulou, 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, Mark Yatskar. *Alexa Prize TaskBot Challenge Proceedings*, 2022.
12. **Visual Goal-Step Inference using wikiHow**
    Yue Yang, **Artemis Panagopoulou**, Qing Lyu, Li Zhang, Mark Yatskar, Chris Callison-Burch. *EMNLP*, 2021. **Oral presentation.**
13. **Induce, Edit, Retrieve: Language Grounded Multimodal Schema for Instructional Video Retrieval**
    Yue Yang, Joongwon Kim, **Artemis Panagopoulou**, Mark Yatskar, Chris Callison-Burch. *arXiv:2111.09276*, 2021.
14. **Self-supervised Optical Flow with Spiking Neural Networks and Event Based Cameras**
    Kenneth Chaney, **Artemis Panagopoulou**, Chankyu Lee, Kaushik Roy, Kostas Daniilidis. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.

Patent Applications and Technical Disclosure

- **Systems and methods for multi-modal language models.** Co-inventor. U.S. application 18/400,477; publication US20240370718A1 (2024).
- **Systems and methods for training and evaluating multimodal neural network based language models.** Co-inventor. U.S. application 18/973,803; publication US20250378347A1 (2025).
- **Systems and methods for visual programming.** Co-inventor. U.S. application 18/972,463; publication US20260044436A1 (2026).
- **Semantic Navigation Using Generative Models and Graphical Maps.** Technical Disclosure Commons, Defensive Publications Series, no. 9308, February 10, 2026. *Technical disclosure.*

Awards and Funding

Best Paper Award, CMCL 2024
Evaluating Vision-Language Models on Bistable Images

CTL Graduate Fellowship for Teaching Excellence, Penn 2023–2024

AWS Research Funding for Fair and Trustworthy AI May 2023

Alexa Prize TaskBot Challenge Finalist 2022

President's Engagement Prize, Penn May 2020

Dean's List, Penn Aug 2017–May 2020

Penn Engineering Exceptional Service Award Mar 2019

Selected Invited Talks and Workshops

Bridging Language and Vision: *Feb 2025*
A Hands-On Introduction to Vision-Language Models
Women in Data Science (WiDS) @ Penn, Philadelphia, PA. Workshop leader.

Bridging Vision and Language: Advances and Challenges *Dec 2024*
Princeton University, Princeton, NJ

Advancing Multimodal AI: Integrating Modalities, Tackling Complex Challenges, and Enhancing Interpretability Sep 2024
University of Pennsylvania, Philadelphia, PA

Academic Service

Journal reviewer: Transactions on Machine Learning Research (TMLR), 2026.

Conference and workshop reviewer: NeurIPS (2026); COLM (2026); CVPR (2025, 2026); ICLR (2026, 2027); ECCV (2024); ACL (2023); ACL Student Research Workshop (2023, 2024); NAACL Student Research Workshop (2025); COLING (2025).

ACL Rolling Review: February and August 2024; February, May, and October 2025; January and May 2026.

Selected Research Coverage and Writing

- **Teaching AI to Read a Map.** Google Research, February 17, 2026. Coauthored research article with Mohit Goyal, introducing MapTrace and its public dataset.
- **ViUniT research coverage.** MarkTechPost, March 7, 2025. Article by Sana Hassan about the visual unit testing framework.
- **I Spy a Metaphor / HAIVMet research coverage.** AI-SCHOLAR, August 31, 2023. Japanese-language article discussing the framework, dataset, and results.

Leadership and Activities

Association of Alumnae (AofA), Penn May 2023–Aug 2026

Board Member, Advancing Women in Engineering (AWE) Aug 2023–Aug 2025

Team Lead, Alexa Prize TaskBot Challenge (Finalist) Aug 2021–May 2022

Women in Computer Science (WiCS) Jan 2019–May 2019