Kate Sanders

Email: ksande25@jhu.edu | Website: katesanders9.github.io | LinkedIn: kate-sanders-395725146

Education

Johns Hopkins University 2021 – Expected early 2026
Ph.D. in Computer Science | Advisor: Benjamin Van Durme *Baltimore, MD*

Johns Hopkins University 2021 – 2023
M.S.E. in Computer Science (GPA: 3.9/4.0) | Advisor: Benjamin Van Durme *Baltimore, MD*

UC Berkeley 2017 – 2020
B.A. in Computer Science (GPA: 3.9/4.0) | Advisor: Ken Goldberg *Berkeley, CA*

Research Interests

Reasoning systems, multimodal understanding, images and video, knowledge verification and grounding, interpretability, datasets and evaluation, reinforcement learning, event semantics, information retrieval, uncertainty.

Skills

Languages: Python, Bash, Java, MATLAB, HTML/CSS, JavaScript.

Tools: PyTorch, Transformers, HuggingFace, Amazon Web Services, Mechanical Turk, Git, LaTeX.

Research Experience

Amazon Web Services May 2025 – Aug. 2025
Applied Scientist Intern *NYC, New York*

- Researching evaluation and RL training paradigms for reasoning models.
- Mentored by Nathaniel Weir and Sapana Chaudhary.

Human Language Technology Center of Excellence June 2024 – Aug. 2024
Program Co-Lead *Baltimore, MD*

- Co-organized and facilitated the 40+ participant, 10-week SCALE 2024 Summer Research Workshop.
- First author of the research paper serving as the workshop basis [16].
- Workshop produced multiple research papers (including [6], [7], [5], [4], [11]) and an ACL 2025 shared task.

Center for Language and Speech Processing, Johns Hopkins University Aug. 2021 – Present
Ph.D. Researcher *Baltimore, MD*

- Researching transparent reasoning [1], [8], [9], [10], [15], multimodal understanding and retrieval [2], [13], [16], [17], and model evaluation [3], [12], [14].

AUTOLab, UC Berkeley Artificial Intelligence Research Sept. 2018 – May 2021
Undergraduate Researcher *Berkeley, CA*

- Led robot error recovery research and improved DexNet system efficiency by 107% [22].
- Trained LSTMs for time series modeling [19] and co-designed a shelf-searching algorithm [18].
- Collaborated to design and deploy web app for computing robot grasp quality [23].

The Miller Lab, UC Berkeley Molecular & Cell Biology Jan. 2018 – May 2018
Research Assistant *Berkeley, CA*

- Developed statistical modeling software in MATLAB for neuronal analysis research [20], [21].

Service

ACL Workshop Organizer (MAGMaR 2025) 2025 – Present

- Co-organizing one-day ACL 2025 workshop on multimodal RAG.

CLSP Application Support Volunteer 2022 – Present

- Coaching underrepresented CS graduate school applicants on writing strong statements of purpose.

Peer Reviewing

- EMNLP 2025, COLM 2025, ACL 2025*, CVPR 2025*, EMNLP 2024, NAACL 2024, NeurIPS 2023, Instruction Workshop @ NeurIPS 2023, NeurIPS 2022, IROS 2021, CASE 2020*[1]

Teaching

Artificial Agents (Co-teacher) Fall 2024
Johns Hopkins University | CS 601.470 *Baltimore, MD*

- Co-wrote and taught syllabus on contemporary AI agent research.

Introduction to Machine Learning (Head TA) Fall 2020
UC Berkeley | CS 189/289A *Berkeley, CA*

- Designed and executed novel course structures for online teaching.
- Led staff of 20+ machine learning TAs and tutors.
- Recieved the UC Berkeley Outstanding Graduate Student Instructor Award.

Adaptive Instruction Methods in Computer Science Spring 2020
UC Berkeley | CS 370 *Berkeley, CA*

- Trained 60+ UC Berkeley EECS TAs and tutors.
- Ran peer tutoring for UC Berkeley's lower-division EECS classes.
- Implemented and maintained tutor-student matching software using Ruby on Rails.
- Co-developed syllabus, assignments, and exams.

Structure and Interpretation of Computer Programs Spring 2019
UC Berkeley | CS 61A *Berkeley, CA*

- Taught discussion and lab sections, hosted office hours, and proctored exams.

Presentations

Grounding Partially-Defined Events in Multimodal Data Nov. 2024
Workshop on the Future of Event Detection, EMNLP 2024 *Talk*

Takeaways from the SCALE 2024 Workshop on Video-based Event Retrieval Sept. 2024
Center for Language and Speech Processing Seminar *Talk*

A Survey of Video Datasets for Grounded Event Understanding June 2024
3rd Visual Datasets Understanding Workshop, CVPR 2024 *Talk*

Multimodal Entailment Trees for Neuro-Symbolic Video Reasoning May 2024
11th Mid-Atlantic Student Colloquium on Speech, Language and Learning *Poster*

Visual Event Semantics Oct. 2023
Center for Language and Speech Processing Seminar *Talk*

[1] *Secondary reviewer

Publications

Google Scholar ID: VJFrfM0AAAAJ

[1] **Sanders, K.**, Van Durme, B. Bonsai: Interpretable Tree-Adaptive Grounded Reasoning. 2025 arXiv preprint.

[2] Martin, A., Kriz, R., Walden, W., **Sanders, K.**, Recknor H., Yang, E., Ferraro, F., Van Durme, B. WikiVideo: Article Generation from Multiple Videos. 2025 arXiv preprint.

[3] Ou, J.*, Walden, W.*, **Sanders, K.**, Jiang, Z., Sun, K., Cheng, J., ..., Van Durme, B. CLAIMCHECK: How Grounded are LLM Critiques of Scientific Papers? 2025 arXiv preprint.

[4] Kriz, R.*, **Sanders, K.***, Etter, D., Murray, M., Carpenter, C., Recknor, H., Blasco, J., Martin, A., Yang, E., Van Durme, B. MultiVENT 2.0: A Massive Multilingual Benchmark for Event-Centric Video Retrieval. CVPR 2025.

[5] Reddy, A., Martin, A., Yang, E., Yates, A., **Sanders, K.**, Murray, K., Kriz, R., M de Melo, C., Van Durme, B., Chellappa, R. Video-ColBERT: Contextualized Late Interaction for Text-to-Video Retrieval. CVPR 2025.

[6] DeGenaro, D., Yang, E., Etter, D., Carpenter, C., **Sanders, K.**, ..., Kriz, R. FORTIFY: Generative Model Fine-tuning with ORPO for ReTrieval Expansion of InFormal NoisY Text ACL 2025 Workshops.

[7] Samuel, S., DeGenaro, D., Guallar-Blasco, J., **Sanders, K.**, Eisape, O., ..., Kriz, R. MMMORRF: Multimodal Multilingual MOdularized Reciprocal Rank Fusion. SIGIR 2025 Demo.

[8] Gupta, K., **Sanders, K.**, Solar-Lezama, A. Randomly Sampled Language Reasoning Problems Reveal Limits of LLMs. ICLR 2025 Workshops.

[9] **Sanders, K.**, Weir, N., Van Durme, B. TV-TREES: Multimodal Entailment Trees for Neuro-Symbolic Video Reasoning. EMNLP 2024.

[10] Weir, N., **Sanders, K.**, Weller, O., Sharma, S., Jiang, D., Jiang, Z., ..., Van Durme, B. Enhancing Systematic Decompositional Natural Language Inference Using Informal Logic. EMNLP 2024.

[11] **Sanders, K.***, Kriz, R.*, Etter, D.*, Recknor, H., Martin, A., Carpenter, C., Lin, J., Van Durme, B. Grounding Partially-Defined Events in Multimodal Data. EMNLP 2024 Findings.

[12] Jiang, Z., Zhang, J., Weir, N., Ebner, S., Wanner, M., **Sanders, K.**, Khashabi, D., Liu, A., Van Durme, B. (2024). Core: Robust Factual Precision Scoring with Informative Sub-Claim Identification. ACL 2025 Findings.

[13] **Sanders, K.**, Van Durme, B. (2024). A Survey of Video Datasets for Grounded Event Understanding. CVPR 2024 Workshops.

[14] Mayfield, J., Yang, E., Lawrie, D., MacAvaney, S., McNamee, P., Oard, D. W., ..., **Sanders, K.**, Mason, M., Hibbler, N. On the Evaluation of Machine-Generated Reports. SIGIR 2024.

[15] Xu, K., Kordi, Y., Nayak, T., Asija, A., Wang, Y., **Sanders, K.**, Byerly, A., Zhang, J., Van Durme, B., Khashabi, D. Tur[k]ingBench: A Challenge Benchmark for Web Agents. NAACL 2025.

[16] **Sanders, K.***, Etter, D.*, Kriz, R.*, Van Durme, B. MultiVENT: Multilingual Videos of Events with Aligned Natural Text. NeurIPS 2023 D&B.

[17] **Sanders, K.**, Kriz, R., Liu, A., Van Durme, B. Ambiguous Images With Human Judgments for Robust Visual Event Classification. NeurIPS 2022 D&B.

[18] Huang, H.*, Dominguez-Kuhne, M.*, Ichnowski, J., Danielczuk, M., Satish, V., **Sanders, K.**, M., Lee, A., Angelova, A., Vanhoucke, V., Goldberg, K. Mechanical Search on Shelves using Lateral Access X-RAY. IROS 2021.

[19] Huh, T. M., **Sanders, K.**, Danielczuk, M., Li, M., Chen, Y., Goldberg, K., Stuart, H. S. A Multi-Chamber Smart Suction Cup for Adaptive Gripping and Haptic Exploration. IROS 2021.

[20] Walker, A., Raliski, B., Nguyen, D., Zhang, P., **Sanders, K.**, Karbasi, K., Miller, E. Imaging Voltage in Complete Neuronal Networks Within Patterned Microislands Reveals Preferential Wiring of Excitatory Hippocampal Neurons. Frontiers in Neuroscience 2021.

[21] Walker, A., Raliski, B., Karbasi, K., Zhang, P., **Sanders, K.**, Miller, E. Optical Spike Detection and Connectivity Analysis With a Far-Red Voltage-Sensitive Fluorophore Reveals Changes to Network Connectivity in Development and Disease. Frontiers in Neuroscience 2021.

[22] **Sanders, K.**, Danielczuk, M., Mahler, J., Tanwani, A., Goldberg, K. Non-Markov Policies to Reduce Sequential Failures in Robot Bin Picking. CASE 2020.

[23] Song, J., Tanwani, A., Ichnowski, J., Danielczuk, M., **Sanders, K.**, Chui, J., Aparicio Ojea, J., Goldberg, K. Robust Task-Directed Grasp Planning as a Service. CASE 2020.