

# Sasha Sax

Research Scientist, FAIR (Meta)

✉ [alexanderesax@gmail.com](mailto:alexanderesax@gmail.com)

🌐 [alex.sax.github.io](https://alex.sax.github.io)

<b>Education</b>	Ph.D. Computer Science [ <i>Maj. AI, Min. Theoretical Statistics</i> ]	2023
	University of California, Berkeley; Berkeley, CA	Advisors: Jitendra Malik & Amir Zamir
	M.S. Computer Science [ <i>Distinction in Research, GPA 4.0+</i> ]	2018
	Stanford University; Stanford, CA	Advisor: Silvio Savarese
B.S. Mathematics [ <i>GPA maj. 3.8</i> ]	2018	
Stanford University; Stanford, CA		
Miscellaneous: AMS Math in Moscow ( <i>Independent University of Moscow/HSE</i> ), Sensory Ecology ( <i>Lund University, Sweden</i> ), Concurrent enrollment ( <i>University of Maryland 2011-12</i> )		

<b>Recent Experience</b>	Computer Vision Research Scientist, FAIR, Meta, Menlo Park, CA	2023-present
	Research Intern, FAIR, Meta; San Francisco, CA	2022-2023
	Research Intern, FAIR, Facebook; Menlo Park, CA	2017-2018
	Software Engineering Intern, Powerpoint, Microsoft,; Mountain View, CA	2016

<b>Awards</b>	Best Paper Award Nomination, CVPR	2020
	Robust Learning Through Cross-Task Consistency	
	Best Paper Award, CVPR	2018
	Taskonomy: Disentangling Task Transfer Learning	
	NVIDIA Pioneering Research Award	2018
	Embodied Real-World Active Perception	
	NVIDIA Graduate Fellowship (Honorable Mention)	2021-2022
	Mid-Level Representations for Robotic Perception	
Stanford University Distinction in Research	2018	
Computational Evidence for Structure in the Space of Tasks		
Winner of CVPR 2019 Habitat Embodied Agents Challenge [RGB Track]	2019	
Mid-Level Visual Representations Improve Generalization and Sample Complexity		
Outstanding Reviewer, CVPR 2020	2020	

<b>Teaching</b>	Machine Learning (TA): <i>Berkeley CS 189/289A</i>	2020
	Representation Learning (Head TA): <i>Stanford CS 331B</i>	2018
	Mathematical Foundations of Computing (TA): <i>Stanford CS 103</i>	2015

<b>Selected Papers</b>	<i>OpenEQA: Embodied Question Answering in the Era of Foundation Models</i>
	<i>Arjun Majumdar*, Anurag Ajay*, Xiaohan Zhang*, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mccvay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent Berges, Shiqi Zhang, Pulkit Agrawal, Yonatan Bisk, Dhruv Batra, Mrinal Kalakrishnan, Franziska Meier, Chris Paxton, <b>Alexander Sax</b>, Aravind Rajeswaran. CVPR, 2024</i>

Alexander Sax, Jeffrey O. Zhang, Bradley Emi, Amir Zamir, Leonidas Guibas, Silvio Savarese  
Jitendra Malik. *CoRL*, 2019. *BayLearn*, 2019.

*Taskonomy: Disentangling Task Transfer Learning* **[Best Paper Award]**

Amir Zamir, Alexander Sax\*, William B. Shen\*, Leonidas Guibas, Jitendra Malik, Silvio Savarese. CVPR, 2018.

*Mid-Level Visual Representations Improve Generalization and Sample Efficiency for Learning Visuomotor Policies*

Alexander Sax, Jeffrey O. Zhang, Bradley Emi, Amir Zamir, Leonidas Guibas, Silvio Savarese  
Jitendra Malik. *CoRL*, 2019. *BayLearn*, 2019.

*Omnidata: A Scalable Pipeline for Multi-Task Mid-Level Vision Datasets from 3D Scans*

Ainaz Eftekhari\*, Alexander Sax\*, Roman Bachmann, Jitendra Malik, Amir Zamir. ICCV, 2021.

*Robust Learning Through Cross-Task Consistency* **[Best Paper Award Nominee, Oral]**

Amir Zamir\*, Alexander Sax\*, Teresa Yeo, Oguzhan Fatih Kar, Nikhil Cheerla, Rohan Suri,  
Zhangjie Cao, Jitendra Malik, Leonidas Guibas. CVPR, 2020.

*Gibson Env: Real-World Perception for Embodied Agents* **[Spotlight]**

Fei Xia\*, Zhiyang He\*, Amir Zamir\*, Alexander Sax, Silvio Savarese. CVPR, 2018.

## Invited Talks

Learning-Based Computational Models of Visual Behavior (Poster), Lund, SE Sep-Oct 2022

This sensory ecology seminar was simply delightful—more embodied AI people should go!!  
You need to **apply early!** It's a 2-week lecture series for 40 PhD/postdocs, taught by the  
world's leading sensory ecologists every 2 years. And you will learn how animals sense and  
represent their environment, and that the line between sensing and thinking is a blurry one.

Mid-Level Visual Representations, CS 280 (guest lecture), Berkeley, CA 2021

Visual Biases in Embodied Agents, Facebook AI Research, Menlo Park, CA Apr. 2019

## Service

**Graduate Mentor:** BAIR Undergraduate Mentoring 2019-2023

**Reviewer:** (CVPR ECCV ICCV ICRA CoRL NeurIPS ICML TPAMI ...) 2018-Present

**Graduate Admissions:** Student Committee 2019, 2021

**Student Organizer:** 3DV Conference 2016

**Junior Class President:** Stanford Class of 2016 2014-2015

**K-5 Math Tutor:** East Palo Alto Tutoring and Tennis 2013-2014

## Demos

Robust Cross-Task Consistency, ECCV 2020, Glasgow. 2020

## Older Experience

**Stanford University, Research Assistant;** Stanford, CA 2016-2017

Visual task relatedness (Taskonomy), Gibson environment. (Silvio Savarese group)

**Microsoft Corporation,** Mountain View, CA Software Engineering Intern, 2016

Powerpoint Designer: improved response time through parallelization (C#), and prototyped a  
logo detector, one of the first internal ML elements in Designer (TF ResNet backbone).

**Stanford University**, Stanford, CA Research Assistant, 2015

Fast square-finding in graphs and fast finding of low-cost replacement paths in presence of edge failures (Virginia Williams group)

**RTI International**, Washington, DC Software Engineering Intern, 2014

Created STATA package to automate statistical analysis and survey ingestion for Early Grade Reading + Math (EGRMA) evaluations in developing countries. Correctly handles reweighting + variance adjustments for multi-level stratified cluster samples. Later used by government orgs in Kenya, Ghana, and Zambia.

**Blackboard Inc.**, San Francisco, CA Software Engineering Intern, 2013

Created an early-warning analytics system to monitor app API traffic, health metrics in real-time. The system used NodeJS, MongoDB, and Hadoop.

**RTI International**, Washington, DC Software Engineering Intern, 2010-2012

Using STATA, developed an automated data-cleaning process that reduced turnaround from 2 months to 1 week and freed PhD statisticians to work on other problems.

## Publications

[11] *OpenEQA: Embodied Question Answering in the Era of Foundation Models*

Arjun Majumdar\*, Anurag Ajay\*, Xiaohan Zhang\*, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mcvay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent Berges, Shiqi Zhang, Pulkit Agrawal, Yonatan Bisk, Dhruv Batra, Mrinal Kalakrishnan, Franziska Meier, Chris Paxton, Alexander Sax, Aravind Rajeswaran. CVPR, 2024

[10] *Omnidata: A Scalable Pipeline for Multi-Task Mid-Level Vision Datasets from 3D Scans*

Ainaz Eftekhari\*, Alexander Sax\*, Roman Bachmann, Jitendra Malik, Amir Zamir. ICCV, 2021.

[9] *Robustness via Cross-Domain Ensembles* **[Oral]**

Teresa Yeo\*, Oğuzhan Fatih Kar\*, Alexander Sax, Amir Zamir. ICCV 2021.

[8] *Robust Policies via Mid-Level Visual Representations: An experimental study in navigation and manipulation*

Bryan Chen\*, Alexander Sax\*, Francis E. Lewis, Silvio Savarese, Jitendra Malik, Amir Zamir, Lerrel Pinto. CoRL, 2020.

[7] *Robust Learning Through Cross-Task Consistency* **[Best Paper Award Nominee, Oral]**

Amir Zamir\*, Alexander Sax\*, Teresa Yeo, Oğuzhan Fatih Kar, Nikhil Cheerla, Rohan Suri, Zhangjie Cao, Jitendra Malik, Leonidas Guibas. CVPR, 2020.

[6] *Side-Tuning: A Baseline for Network Adaptation via Additive Side Networks* **[Spotlight]**

Jeffrey O. Zhang, Alexander Sax, Amir Zamir, Leonidas Guibas, Jitendra Malik. ECCV, 2020.

[5] *Learning to Navigate via Mid-Level Visual Priors*

Alexander Sax, Jeffrey O. Zhang, Bradley Emi, Amir Zamir, Leonidas Guibas, Silvio Savarese, Jitendra Malik. CoRL, 2019.

[4] *Mid-Level Visual Representations Improve Generalization and Sample Efficiency for Learning Visuomotor Policies* **[Oral]**

Alexander Sax, Jeffrey O. Zhang, Bradley Emi, Amir Zamir, Leonidas Guibas, Silvio Savarese Jitendra Malik. *Arxiv* 2018. *BayLearn*, 2019. (Oral)

[3] *Taskonomy: Disentangling Task Transfer Learning* **[Best Paper Award, Oral]**

Amir Zamir, Alexander Sax\*, William B. Shen\*, Leonidas Guibas, Jitendra Malik, Silvio Savarese. *CVPR*, 2018.

[2] *Embodied Real-World Active Perception* **[Spotlight]**

Fei Xia\*, Zhiyang He\*, Amir Zamir\*, Alexander Sax, Silvio Savarese. *CVPR*, 2018. (Spotlight)

[1] *Joint 2D-3D-Semantic Data for Indoor Scene Understanding*

Iro Armeni\*, Alexander Sax\*, Amir Zamir\*, Silvio Savarese. *Arxiv (preprint)*, 2016.