

Maximilian Du

maxjdu@stanford.edu / maximiliandu.com / (315) 751-1504

OBJECTIVE

To obtain a graduate research position in the field of computer science. Interested in creating broadly intelligent robots and other agents that learn from diverse data, actively perceive its environment, and adapt through real-world interaction. Qualified with three years of research experience in reinforcement learning, imitation learning, computer vision, and real robots. Further academic background in mathematics, creative/technical writing, and psychology.

EDUCATION

Stanford University	09/2020 – Present
BS, Computer Science (AI Track) Minor, Creative Writing Minor, Psychology	
GPA: 4.110 / 4.0. Technical GPA: 4.092 / 4.0	
Fayetteville-Manlius High School	09/2016 – 06/2020
High School Diploma	
GPA: 103 / 100	

RESEARCH EXPERIENCE

Undergraduate Research Assistant	01/2021 – Present
Stanford Artificial Intelligence Laboratory (SAIL) IRIS Lab	
<ul style="list-style-type: none"> Developed robot learning methods that use multiple sensory modalities, learn from diverse data, and adapt quickly to new tasks. Led two projects and co-led a third. Roles included proposing project direction, initiating experiments, analyzing results, presenting progress at weekly meetings, and creating paper figures & presentations Working under Prof. Chelsea Finn. Collaborated with Profs. Dorsa Sadigh, and Tobias Gerstenberg. Mentored by Suraj Nair and Alexander Khazatsky Published two papers in top robotics conferences: Robotics: Science and Systems 2022 & 2023. Presented papers in reading group and hosted summer group meetings for undergraduates 	
When at First You Don't Succeed: Knowing When to Try Again in Novel Test-Time Scenarios	02/2023 – Present
<i>Current work. Advised by Sasha Khazatsky, Chelsea Finn & Tobias Gerstenberg</i>	
<ul style="list-style-type: none"> Making a robot-learning method that adapts quickly from environment interactions by learning a repertoire of strategies and searching for the best strategy during a test-time task Designed and implemented a counterfactual reasoning model that compares the robot to a hypothetical expert, allowing the robot to decide when to try another strategy Designed visualizations for various parts of the algorithm, allowing for rapid diagnosis of model behavior 	
Behavior Retrieval: Few-Shot Imitation Learning by Querying Unlabeled Datasets	02/2022 – 01/2023
<i>Published at Robotics: Science and Systems 2023. Advised by Suraj Nair, Chelsea Finn & Dorsa Sadigh</i>	
<ul style="list-style-type: none"> Proposed an approach to reuse previously collected datasets of robot interactions by filtering data according to its relevance to a target task Designed, implemented, and trained an encoder that maps dataset points to a latent space. Distance in the latent space allows the separation of relevant and irrelevant data Ran hundreds of experiments in MuJoCo simulation to make experimentally-driven design choices in the algorithm Proposed and carried out experiments on a real Widowx robot arm, amounting to over 1000 evaluation trials and three environments, showing that the filtered data could boost behavior cloning performance by over 35% 	
Play it by Ear: Learning Skills amidst Occlusion through Audio-Visual Imitation Learning	01/2021 – 01/2022
<i>Published at Robotics: Science and Systems 2022. Advised by Suraj Nair & Chelsea Finn.</i>	

- Created a robot-learning method that leverages images and interaction sounds to improve decision-making through signals from multiple sensory modalities.
- Implemented, and tested a collection of reinforcement learning & behavior cloning approaches, leading to a data-driven choice of the final algorithm
- Proposed and iterated on difficult visually-occluded tasks for a real and simulated robot arm. After training and fine-tuning on expert supervision, the robot could find and extract keys from an opaque bag up to 70% of the time

Improving LSTM Neural Networks for Better Short-Term Wind Power Predictions 06/2018 – 11/2019

Published at IEEE Renewable Energy and Power Engineering 2019. Advised by Joshua Comden & Zhenhua Liu

- Created an autoregressive model that predicts wind power outputs, allowing for greater stability of power grids
- Proposed new metrics to evaluate power prediction models and allow for targeted improvements
- Collected, processed, and combined large (3TB) databases of wind power and weather forecast data, which provided important context for the models and led to improvements in accuracy

PUBLICATIONS

- **Maximilian Du**, Sasha Khazatsky, Tobias Gerstenberg, and Chelsea Finn (2024). When at First You Don't Succeed: Knowing When to Try Again in Novel Test-Time Scenarios. *In Preparation for Robotics: Science and Systems XX*
- **Maximilian Du**, Suraj Nair, Dorsa Sadigh, and Chelsea Finn (2023). Behavior Retrieval: Few-Shot Imitation Learning by Querying Unlabeled Datasets. *Robotics: Science and Systems XVIV*
- **Maximilian Du***, Olivia Lee*, Suraj Nair, and Chelsea Finn (2022). Play It by Ear: Learning Skills amidst Occlusion through Audio-Visual Imitation Learning. *Robotics: Science and Systems XVIII*
- **Maximilian Du** (2019). Improving LSTM Neural Networks for Better Short-Term Wind Power Predictions. *IEEE 2nd International Conference on Renewable Energy and Power Engineering*
- Homer Walke, Kevin Black, Abraham Lee, Moo Jin Kim, **Maximilian Du**, Chongyi Zheng, Tony Zhao, Philippe Hansen-Estruch, Quan Vuong, Andre He, Vivek Myers, Kuan Fang, Chelsea Finn, and Sergey Levine (2023). BridgeData V2: A Dataset for Robot Learning at Scale. *Conference on Robot Learning (CoRL)*
- Open-X-Embodiment Collaboration (2023). Open X-Embodiment: Robotic Learning Datasets and RT-X Models. *LangRob Workshop, CoRL 2023*

RELEVANT COURSEWORK

Graduate Computer Science: CS 223A (Robotics), CS229M (ML Theory), CS 236 (Deep Generative Models), CS 224R (Deep RL), CS 234 (RL), CS 224N (Deep learning for NLP), CS 330 (Deep Multi-task and Meta Learning), CS 231N (Deep Learning for Computer Vision), CS 229 (ML), UC Berkeley CS 285 (Deep RL, self-study)

Undergraduate Computer Science: CS 161 (Algorithms), CS 110 (Computer Systems Principles), CS 107E (Introductory Computer Systems), CS106B (Programming Abstractions)

Mathematics: EE 364A (Convex Optimization), EE 276 (Information Theory), Math 115 (Real Analysis), CS 228 (Probabilistic Graphical Models), Math 113 (Linear Algebra and Matrix Theory), Math 51 (Linear Algebra and Multivariable Calculus)

Psychology: Psych 226 (Models and Mechanisms of Memory), Psych 169 (Advanced Seminar on Memory), Psych 45 (Learning & Memory), Psych 30 (Perception), Psych 50 (Cognitive Neuroscience)

Writing & Humanities: English 290 (Advanced Fiction), English 191 (Intermediate Non-Fiction), English 190 (Intermediate Fiction), English 92 (Introductory Poetry), English 91 (Introductory Creative Non-Fiction), Phil 2 (Moral Philosophy)

HONORS AND AWARDS

Hertz Fellowship Finalist. Top 45 out of 860 applicants	01/2024
CRA Outstanding Undergraduate Researcher Award Finalist	01/2024
Stanford CS Honors Program	09/2023 – Present
Stanford Undergraduate Creative Writing Prize. Third place prize out of 100+ submissions.	05/2023
Stanford Small Grant Recipient. Faculty-endorsed research in humanities.	03/2023 – 09/2023
Tau Beta Pi Engineering Honor Society Candidate. Top 12.5% in School of Engineering	01/2023 – Present

Stanford CS231N Final Project Winner. Selected out of 370+ students	06/2022
Lunsford Award Finalist For Oral Presentation of Research. Top 5 out of 600+ class presentations	03/2022
Stanford CS109 Final Project Winner. Selected out of 200+ students	06/2021
Regeneron Science Talent Search Semifinalist. Top 300 applicant from 1900+ submissions	01/2020

TEACHING AND OUTREACH

Reviewer	10/2023
Deployable Robot Learning Workshop at Conference on Robot Learning (CoRL) 2023	
Deep Learning Portal Mentor	Starting 01/2024
Stanford Computer Science Department	
<ul style="list-style-type: none"> Will help disadvantaged students learn AI by hosting live, weekly office hours 	
Stanford Splash Lecturer	11/2021 – Present
Stanford Educational Studies Program	
<ul style="list-style-type: none"> Designed a curriculum to introduce high school students to robot learning through connections to animal training and other fields of psychology Taught five course iterations so far, with roughly 200 total students. Received numerous positive student reviews. Invited to present similar lecture at the International Marine Mammal Trainers Association (IMATA) conference in 2024 	
CS 106A/B Section Leader (TA)	01/2021 – 06/2022
Stanford Computer Science Department	
<ul style="list-style-type: none"> Led weekly small-group sections for the popular CS106A/B Stanford course series. Helped with conceptual and coding problems during weekly office hours, graded homework assignments and exams Received crowd-voted office hour service award and positive student feedback through anonymous mid-quarter evaluation. 	

LEADERSHIP

Senior Producer & Writer	06/2022 – Present
Stanford Storytelling Project	
<ul style="list-style-type: none"> Leading production team on a podcast series, to be published on PRX, PodBean, and aired on KZSU radio Doing fieldwork, interviews (30+ hours) and archival research (5k+ documents) for a creative nonfiction book on the human-animal relationship. Work is partially funded by Stanford University. Presented findings on the human-animal relationship at Stanford Symposium of Undergraduate Research and Public Service (SURPS) 	
Volunteer Advisor + Narrative Lead	03/2023 – 06/2023
Truth4Toki Advocacy Group	
<ul style="list-style-type: none"> Helped organize 25 zoological professionals to advocate for animal welfare by creating a unified narrative, editing written posts, and preparing people for media interviews Helped gain 40k+ signatures on a petition and prolonged local & national media attention (Good Morning America, NBC Seattle, WPLG Miami, VICE) 	

SKILLS

Languages: Python, C++, C, Java, HTML/CSS

ML & Data Tools: PyTorch, Tensorflow, Numpy, Matplotlib, Linux, Git, Zotero

Codebase/API Familiarity: Robosuite, Robomimic, Roboverse, MuJoCo, PyBullet, SLURM, ROS

Robots: Franka-Emika Panda Arm, Widowx Arm,

Other Tools: L^AT_EX, Adobe Illustrator / Audition / Premiere Pro, CAD Design, Oscilloscope, Soldering

Other Skills: Archival Document Organization, Narrative Interviews, Audio Production, Educational Presentations, Narrative Theory & Storytelling

OTHER PROJECTS

The Basics: Research Toolkit

12/2021 – Present

Personal Project

- Kept track of code snippets commonly used during research and combined them into a public GitHub repository
- Curated Zotero database of 800+ research papers with meaningful organization scheme
- Wrote hundreds of pages of notes on math and AI subjects that are uploaded to my website

Archivist & Editing Work

05/2023 – Present

Themed Reality Journalism Group

- Trained people to use archival software for FOIAs and AI tools for content summary
- Edited written articles and scripts related to the zoological industry, leading to multi-million video views and high-influence op-eds

Policy Evaluation for Berkeley and Google

01/2023 – 09/2023

Stanford IRIS Lab

- Assisted with research projects at Google and Berkeley by evaluating policies on a Widowx robot
- Assisted with some data analysis and evaluation task proposal, leading to results presented in published papers

Looking Under the Hood of DetectGPT

01/2023 – 03/2023

CS 224N Final Project

- Extended published results on DetectGPT, an algorithm that detects large language model output
- Proposed and tested ways of improving DetectGPT by focusing on certain parts of speech

Can you Macgyver It? Teaching an Agent to Use Tools

01/2023 – 03/2023

CS 234 Final Project

- Implemented a policy gradient algorithm to solve a tool-usage environment
- Explored impacts of different exploration algorithms on data efficiency and final performance

Sixteen Pixels is (Almost) All You Need: Crafting Parameterized Image Uncrumpling Models

03/2022 – 06/2022

CS 231N Final Project | Best project winner

- Modified the Pix2Pix algorithm to take in a crumpled image and output its uncrumpled form
- Proposed a smaller PatchGAN architecture that qualitatively outperforms existing PatchGAN architectures

MidiStyle: Parameterized Audio Style Transfer for Instrument Swapping

09/2021 – 11/2021

CS 229 Final Project

- Used a convolutional autoencoder to transform piano music into other instruments
- Used a FiLM-style conditioning to modify the output instrument