

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 robots and other intelligent agents that get closer to human-like intelligence by learning from diverse data, leveraging multimodal sensory inputs, and adapting 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: 4.108 / 4.0

Fayetteville-Manlius High School **09/2016 – 06/2020**
High School Diploma
GPA: 103 / 100

PUBLICATIONS

- 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 et al. (2023). BridgeData V2: A Dataset for Robot Learning at Scale. *Conference on Robot Learning (CoRL)*

RESEARCH EXPERIENCE

Undergraduate Research Assistant **01/2021 – Present**
Stanford Artificial Intelligence Laboratory (SAIL) | IRIS Lab

- Developed robot learning algorithms that use multiple streams of information, learn from relevant demonstrations, and retry meaningfully after failed attempts.
- Led multiple projects, including proposing project direction, initiating experiments, analyzing results, presenting progress at weekly meetings, and creating paper figures & presentations
- Published two papers in top robotics conferences: RSS 2022, and RSS 2023.
- Advised by Prof. Chelsea Finn. Collaborated with Profs. Dorsa Sadigh, and Tobias Gerstenberg
- Presented papers in reading group and hosted summer group meetings for undergraduates

Try, Try Again: Behavior Cloning for Novel Test-Time Scenarios **02/2023 – Present**
Advised by Sasha Khazatsky, Chelsea Finn & Tobias Gerstenberg

- Creating an algorithm that allows robots to try new strategies upon task failure, leading to higher success rates
- Designed and implemented a progress-measuring model, 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

Advised by Suraj Nair, Chelsea Finn & Dorsa Sadigh

- Created a novel data selection algorithm that reused previously collected datasets of robot interactions, boosting success rates from behavior cloning by over 35%
- Designed and implemented a state-similarity classifier that allowed 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, verifying the properties of our method

Play it by Ear: Learning Skills amidst Occlusion through Audio-Visual Imitation Learning

01/2021 – 01/2022

Advised by Suraj Nair & Chelsea Finn. Co-led with Olivia Lee.

- Created an end-to-end robot learning algorithm that fused audio, visual, and proprioceptive data, improving success rates of hard tasks by over 50% percent.
- Proposed, 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

Improving LSTM Neural Networks for Better Short-Term Wind Power Predictions

06/2018 – 11/2019

Advised by Joshua Comden & Zhenhua Liu

- Created a sequence model that uses weather data to accurately predict short-term wind power outputs
- Proposed new metrics to measure the performance of time series models
- Collected, processed, and combined large (3TB) databases of wind power and weather forecast data

RELEVANT COURSEWORK

Graduate Computer Science: 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 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 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)

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

HONORS AND AWARDS

CRA Outstanding Undergraduate Researcher Award Nomination, Fall 2023. One of four undergraduates nominated by the Stanford CS department. Currently applying for the award.

Stanford CS231N Final Project Winner, Spring 2022. Selected out of 370+ students

Stanford CS109 Final Project Winner, Spring 2021. Selected out of 200+ students

Stanford Undergraduate Creative Writing Prize, 2023. Third place prize out of 500+ submissions

Stanford Small Grant Recipient, 2023

Stanford Tau Beta Pi Engineering Honor Society, Candidate Fall 2022

Stanford Lunsford Award For Oral Presentation of Research, 2022 Finalist

Regeneron Science Talent Search, 2020 Semifinalist (Scholar)

Discovery Education “Making For Good” Challenge, 2020 National Second Place

TEACHING AND OUTREACH

Reviewer

Deployable Robot Learning Workshop at Conference on Robot Learning (CoRL) 2023

Deep Learning Portal Mentor

Starting 01/2024

Stanford Computer Science Department

- Will help disadvantaged students learn AI by hosting live, weekly office hours
- Will explain difficult concepts, debug machine learning code, and find resources for students

Stanford Splash Lecturer

11/2021 – Present

Stanford Educational Studies Program

- Designed a curriculum from scratch to introduce high school students to reinforcement learning by using connections to animal training and other fields of psychology
- Iterated and taught over four course cycles, with roughly 200 total students

CS 106A/B Section Leader (TA)

06/2021 – 06/2022

Stanford Computer Science Department

- Led weekly small-group sections for the popular CS106A/B Stanford course series.
- Helped with conceptual and coding problems during weekly office hour shift
- Graded homework assignments and course exams

LEADERSHIP

Senior Producer & Writer

06/2022 – Present

Stanford Storytelling Project

- Leading production team on a one-hour audio story, 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

Volunteer Advisor + Narrative Lead

03/2023 – 06/2023

Truth4Toki Advocacy Group

- 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 prolonged local and national media attention (Good Morning America, NBC Seattle, WPLG Miami, VICE)
- Ran organization website and helped gain more than 40k signatures on Change.org

OTHER PROJECTS

- Policy Evaluation for Berkeley and Google** **01/2023 – Present**
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
- 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