

Joseph P. Robinson

ML ENGINEER · RESEARCHER · ENTREPRENEUR

+1 (978) 764 0003 | jobs.vision@gmail.com | jobsvision.com | github.com/jvision | linkedin.com/in/visionjo | [M jvision](https://m.jvision) | scholar

“Be like water making its way through cracks. Do not be assertive, but adjust to the object, and you shall find a way around or through it. If nothing within you stays rigid, outward things will disclose themselves.” –Bruce Lee

Experience

BitHuman.io | *Head of Rapid Product Integration and Development (RAPID)* | Remote Aug 2023 - Jul 2024

- Managed five SW and two machine learning engineers to build MVP: a human-like AI agent tailored to client brands.
- Built the next generation product—an AI agent via a fully generative multimodal solution optimized to run on-prem, on the cloud, or hybrid: ridding the dependency of computationally demanding graphic rendering and designed to scale.
- Optimized system using start-of-the-art components (*i.e.*, talking face, TTS, SST, image correction/super-resolution, and lightweight and global LLMs) and data (*i.e.*, tensors) passing for the multimodal solution.
- Bridged gap between product, engineering, research, sales, and creative with weekly deep dives, product requirements, and backlogs; refined company priorities based on latest tech and customer feedback; led bi-weekly sprints, ensuring rapid production and QA.
- Designed goal-orientated paradigm and a novel multi-modal recommendation schema.
- Filed four patents as part of effort to extend company IP.

Northeastern University | *Machine Learning Performance Engineer* | Remote Jan 2023 - Dec 2023

- Member of [Research Computing](#) supporting ML/AI/HCI research for faculty & researchers on SLURM-based HPC cluster.
- Improved the reliability and efficiency of GPU-based S.W. while optimizing the performance of ML models on our HPC.
- Developed an LLM as a user-facing Q&A system to answer cluster-specific questions using short-term and long-term memory.
- Created a micro-service to recommend optimal Slurm configurations for users running jobs on the cluster.
- Developed [user-facing documentation](#): revamped content, extended scope, improved look-and-feel, and delegated sections.

Tufts University | *Part-Time Lecturer- Machine Learning* | Medford, MA Jan 2022 - Dec 2022

- Taught Machine Learning, project-based course with 120 students and 3 TAs.
- Leveraged [otto-grader](#) to automatically grade and give immediate feedback: implemented extension to process Jupyter NBs.

Vicarious Surgical | *AI Engineer- Team Perception* | Waltham, MA Mar 2021 - July 2022

- Built depth perception system for complex human tissues and organs seen by the robotic arm with surgical precision.
- Deployed real-time depth perception SW in C++ and CUDA for a hand-held device prototype running an [NVIDIA's Jetson](#).
- Collected simulated (*i.e.*, [ROS](#) and [Gazebo](#)) and real data (*i.e.*, actual system) of the robotic arm to train and then fine-tune deep models via multi-task and transfer learning, respectively; also, allowing for precise performance measures.
- Managed team's Confluence: used its spaces organized with macros and templates – used to standardized data collection via a Data Factory Space using template [Datashheets](#) for proper data characterization.
- Developed an end-to-end virtual assistant using [NVIDIA's RIVA](#) to control the surgical robotic via voice commands at the character-level to allow the system to perform best in the medical (*i.e.*, , surgical) domain.

ISMConnect | *Grad Student Intern* | Cambridge, MA May 2019 - Aug 2019

- Acquired a deep understanding of problems of bias in facial recognition technology (*i.e.*, “Fairness in ML”).
- Designed a characterization technique via fundamentals of signal detection theory and built [BFW dataset](#) to measure model bias [4].
- Improved state-of-the-art deep CNNs were obtained across under-represented subgroups. Furthermore, improvements yielded performance ratings consistent across different demographics [1].
- Proposed feature learning scheme to de-bias face features and balance subgroup performances, increasing overall performance.

Snap Inc. (Snapchat) | *Grad Student Intern* | Santa Monica, CA May 2018 - Aug 2018

- Proposed a Laplace KL-divergence objective to train landmark localization models that are more certain about predictions [6].
- Implemented an adversarial training framework that leverages large amounts of unlabeled data during training.
- Designed experiments to show that our model outperforms recent works in face landmark detection, along with ablation studies that reveal our model performs comparably well at 1/8 its original size (*i.e.*, <160 Kb) and in real-time (*i.e.*, >20 fps).

Systems & Technology Research (STR) | *Grad Student Intern* | Woburn, MA May 2016 - Sep 2017

- Worked on [IARPA's JANUS Project](#) (Phase II): Developed C++ API of clustering algorithms (*e.g.*, K-means, GMM, Agglomerative) to generate results on huge facial image datasets reported at NIST data calls. Our team earned top standing per NIST evaluation.
- Implemented Product Quantization & Nearest Neighbor Search via Inverted File Structure to gain a 200x speedup in C++ production code with negligible performance drop.
- Contributed of [IARPA's Odin Program](#) (Phase I): Designed & implemented a Python API for adversary attacks on ML models; designed several benchmarks for adversary attacks using: black-box regime using various deep CNNs for face recognition.

MIT Lincoln Labs (MIT-LL) | *Grad Student Intern* | Lincoln, MA May 2014 - Aug 2014

- Led joint team (*i.e.*, SMILE and MIT-Lincoln Lab) in [TRECVID](#) debut— 3rd highest accuracy in [Multimedia Event Detection](#) (MED'15) via CNNs pre-trained on auxiliary data to encode events as scenes and objects [9].
- Encoded videos with two pre-trained CNNs (*i.e.*, 1,000 objects and 360 scenes), fused features, and trained SVMs.
- Achieved 3rd in MED'15—detect complex events in a large corpus of videos with distractors and hard negatives.

BBN Technology | *DSP Engineering Co-Op* | Cambridge, MA Jan 2013 - Sep 2013

- Worked on the [Helicopter Alert and threat Termination-Acoustic System](#) to detect, localize, and classify small arms from helicopters.
- Built a Java tool to analyze audio data; ultimately, I improved the system from 86% to 92% via mean template matching.

Analogic Cooperation | *EE Image Reconstruction Co-Op* | Peabody, MA Jan 2012 - Sep 2012

- Focused on SW components: critical image processing and detection algorithms in airport bag-scanners to transition from single to dual energy [CT system](#).
- Optimized algorithms with GPGPU and Intel Vector library to reduce HW requirements from 3-to-2 computers.
- Established training protocols and built graphical interface in JAVA for employees to learn how to visually inspect for threats.

Education

Ph.D. in Computer Engineering, Northeastern University (NEU) | Boston, USA

Jan 2015 - Dec 2020

BS in Electrical & Computer Engineering, Northeastern University | Boston, USA

Jan 2011 - May 2014

Achievements: 1st @ College-wide Senior Capstone Competition | 1st @ ECE Dept. - Remote Control Contest | Best Senator (2x).
15+ International Conferences: CVPR, ACM-MM, ICCV, ECCV, AMFG, ICME, AAAI | Hosted workshops and tutorials at numerous top-tier conferences | 3x Best Reviewer for IEEE AMFG | Organized and Hosted New England Computer Vision (NECV) Conference at NEU.
Extra Curriculum: President, IEEE | Lead Research Ambassador, Student Research Engagement Committee | Student Senator, SGA | completed two research experience for undergrads (REU) in 2010 and 2011.
Courses: Deep Learning | NLP | Optimization | Machine Learning | Computer Vision | Digital Image Proc. | Parallel Data Processing.

Selected Publications

30+ peer review papers. See [Scholar](#) or <https://www.jrobs-vision.com/publications>. Also, blogger: <https://medium.com/@jvision>.

- [1] **JP Robinson** et al. “Balancing Biases and Preserving Privacy on Balanced Faces in the Wild”. In: *TIP*. 2023.
- [2] **JP Robinson** et al. “Families In Wild Multimedia: A Multi-Modal Database for Recognizing Kinship”. In: *TMM* (2021).
- [3] **JP Robinson** et al. “Survey on the Analysis and Modeling of Visual Kinship: A Decade in the Making”. In: *TPAMI* (2021).
- [4] **JP Robinson** et al. “Face Recognition: Too Bias, or Not Too Bias?”. In: *CVPR* (2020).
- [5] **JP Robinson** et al. “Recognizing Families In the Wild: The 4th Edition”. In: *FG*. IEEE. 2020.
- [6] **JP Robinson** et al. “Laplace landmark localization”. In: *ICCV* (2019).
- [7] **JP Robinson** et al. “To Recognize Families In the Wild: A Machine Vision Tutorial”. In: *ACM MM*. 2018.
- [8] **JP Robinson** et al. “Visual kinship recognition of families in the wild”. In: *TPAMI* (2018).
- [9] **JP Robinson** and Y Fu. “Pre-trained D-CNN models for detecting complex events in unconstrained videos”. In: *Sensing and Analysis Technologies for Biomedical and Cognitive Applications*. International Society for Optics and Photonics. 2016.
- [10] **JP Robinson** et al. “Families in the wild (fiw): Large-scale kinship image database and benchmarks”. In: *ACM MM*. 2016.

Selected Projects

Visual Kinship Recognition – Families In the Wild (FIW) Dataset

Sep 2015 - Dec 2020

Northeastern University

- Built the largest and most comprehensive dataset for visual kinship recognition – 1,000 families each with 10+ family photos [8].
- Established diverse team for data collection world-wide; provided JAVA GUI with semi-supervised model for rapid labeling [10].
- Generated several benchmarks and novel models and algorithms.
- Extended image set with multimedia (*i.e.*, video and audio) with a proposed multimodal labeling algorithm [2].
- Proposed generative model to map face pair (*i.e.*, prospective parents) to a child – designed to sweep age and gender via latent space.
- Published 12+ peer-reviewed papers, including a comprehensive survey [3], 5-editions of a big data challenge as part of top-tier proceedings (*e.g.*, [5]), and organized and delivered tutorials on the topic at CVPR, AMFG, and ACM-MM (*e.g.*, [7]); Hosted [Kaggle Competition](#) and working to establish proceedings to invite +550 teams to publish work. See thesis on [ArXiv](#).

Tunnel Inspecting Robot

Sep 2013 - May 2014

Northeastern University

- Built a tunnel inspecting robot to address labor-intensive processes that put the inspector and drivers beneath in harm’s way.
- Completed a design with a tight budget of \$300 using Raspberry Pi and Arduino micro-controllers to capture/stream two video views and control base motors/servos, respectively.
- Served as team leader: ensured deliverables met deadlines via sprint-ban, and held daily stand-ups, weekly deep dives, and biweekly sync-ups with our advisor; also, worked closely with MASS-DOT for product requirements.
- Implemented computer vision algorithms to detect and record cracks and measure to determine the rate at of change.
- Developed JAVA GUI for users to log in, record progress, and access the database. Also, implemented the drivers to send signals from an Xbox control programmed to control the robot remotely (*i.e.*, two joysticks to control the base and adjustable arm camera).
- Used TCP for real-time streaming of two videos from Pi to host desktop; configured CV algorithms to run in background on the host.
- Won 1st in Senior Capstone competition based on technology, presentation (oral & poster), and live demos (system & simulation).

Selected Patents

Joseph P Robinson, Yun Fu, Method of Providing Personalized Customer Interactions with Adaptive AI. *U.S. Patent 18,732,610*.

March 2024

Joseph P Robinson et al., Dual-Layered Artificial Intelligence System. *U.S. Patent 18,604,504*.

March 2024

Joseph P Robinson et al., Dual-Layered AI System w Large Language Models and Different Visual Agents. *U.S. Patent 18,634,939*.

April 2024

Joseph P Robinson, Yun Fu, Background Noise Filtering System. *U.S. Patent 18,634,991*.

April 2024

Selected Certificates

- | | | |
|------|--|---------------|
| 2012 | Gordon Engineering Leadership Program , CENSSIS – Northeastern University | Boston, MA |
| 2010 | Computer Aided Drafting Certificate , Northern Essex Community College | Haverhill, MA |
| 2007 | Home Improvement Contractor (HIC) License , The Commonwealth of Massachusetts | Boston, MA |
| 2006 | Massachusetts Construction Supervisor’s License (CSL) , Merrimack College | Merrimack, MA |

Skills

Programming	Python, C/C++, CUDA, Java, Catkin, CMake, Matlab, Git, Scripting (Bash), LaTeX, HTML, Emacs
Software	Linux, Tensorflow, Pytorch, Docker, Kubernetes, OpenCV, Keras, Spark, Slurm, Solidworks, MPI, ROS 1/2, Gazebo
Soft Skills	Time Management, Problem-solving, Documentation, Presenting, Leadership, On-site coordination, Teaching, Writing.