

Cell: +1 978 764 0003
Salisbury, MA
jobs.vision@gmail.com

Joseph P. Robinson, Ph.D.

AI Research Engineer

Web: jobs-vision.com
GitHub: github.com/visionjo
LinkedIn: linkedin.com/in/jrobby

‘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

WORK EXPERIENCE

Vicarious Surgical

AI ENGINEER, TEAM PERCEPTION

Tackled depth perception (*i.e.*, monocular, stereo, and MVS) for 3D reconstruction from images and videos. Completed lit review and presented a four-part series of the SOTA to find the best for our surgical endoscope (*i.e.*, human tissues and organs seen by the robotic arm). Used SOTA deep models and post-processing to estimate the depth and its confidence maps at the target fps and accuracy. Acquired labeled data to report ratings for our robot–data capture required learning ROS and Gazebo. Designed and collected both simulated and *real* datasets to train and fine-tune deep models via multi-task and transfer learning; then precisely measured system performance (*i.e.*, characterize and find shortcomings). The system was the basis of the critical AI components. It was later deployed in another company effort, proving it was effective at generalizing and easily portable. The system was implemented in C++ and CUDA, optimizing it to run on NVidia’s Jetson. Other contributions include, but are not limited to, the design and implementation of a virtual assistant using NVIDIA RIVA for controlling via voice and natural language: built both word and character-level, as the latter proved more extendable and generalized best to the surgical domain. Learned Confluence, managed spaces, and created macros and templates: created a Data Factory Space with a Datasheet template, dataset exploration pages, how-to’s (*e.g.*, camera calibration), and more; also organized weekly paper studies, which matured into a computer vision textbook study, with the schedule (*i.e.*, the chapters and corresponding papers for deep dives) and material per session.

March 2021 — Present

Waltham, MA

ISMConnect

GRAD STUDENT INTERN, VISION GROUP

My task was to measure the fairness of SOTA facial recognition (FR) systems across gender and ethnicity. Built the image-based BFW dataset as a proxy, supported by appropriate evaluation protocols, various metrics, and benchmarks. Fundamental signal processes motivated the novel adaptive thresholds that proved effective in the FR system used in the product, assuming a prior knowledge of the demographics of the subject of interest [3]. We generalized this by learning a mapping from the feature space of the respective model to one that removes demographic information while preserving the knowledge needed to identify ([under review](#)): a means to debias face embeddings while enforcing privacy. Download data, get source code, and learn more on [GitHub](#).

May 2019 — Aug 2019

Cambridge, MA

Snap Inc.

GRAD STUDENT INTERN, RESEARCH TEAM

Implemented a novel SOTA face landmark detection model with reduced size and sped up [4]. Proposed LaplaceKL loss function: loss function that extends the renowned soft-argmax with additional statistics to yield heatmaps of higher confidence (*i.e.*, precision) and, thus, improved accuracy. Trained with few labeled samples and many unlabeled ones, using adversarial learning based on produced heatmaps. Explored reduced sizes - competitive with SOTA for reduced to 170KB running in real-time.

May 2018 — Aug 2018

Santa Monica, CA

STR Technology

GRAD STUDENT INTERN, VISION GROUP

For IARPA’s Odin Program (2017): designed and implemented a Python API for adversary attacks on ML models; assessed adversary attacks (black-box regime) using different SOTA CNNs for face recognition. For IARPA’s JANUS Program (2016): built clustering API in C++ to generate results for NIST data calls; implemented PQ for NN search via IFS: 200x speedup, negligible performance drop.

May 2016 — Sep 2017

Woburn, MA

MIT Lincoln Labs

ENGINEERING INTERN, VIDEO AND IMAGE UNDERSTANDING

Led team (*i.e.*, SMILE-MIT-LL) in TRECVID debut in MED’15 (3rd place). Encoded videos with two CNNs (*i.e.*, 1,000 objects and 360 scenes), fused and trained SVMs to detect complex events in a large video corpus with *distractors* and hard-negatives.

May 2014 — Aug 2014

Lexington, MA

Raytheon BBN Technology

DSP ENGINEERING CO-OP

Worked on HALTT: gun detection and localization for helicopters. Improved system from 86% to 92% via template matching. Created a harness to simulate gunshots for the system by injecting audio signatures of different guns. Built MATLAB tools to analyze detection algorithms for any clip captured during field tests or in system logs. Built JAVA tools for customers to visualize data.

Jan 2013 — Sep 2013

Cambridge, MA

Analogic Corporation

EE IMAGE RECONSTRUCTION CO-OP

Focused on SW components: image processing and detection algorithms in [airport bag scanners](#). Optimized code with GPGPU and Intel Vector library, reducing HW requirements (3-to-2 computers) in transition from single-to dual-energy CT systems. Profiled GPUs, providing incentives for system upgrades. Built ImageJ plugin for threat analysis. Designed a database to automate logging and report generation. Built a sister plugin to train and certify employees in data analysis and proper understanding of possible threats. After completion of training and test modules, they can analyze field results and tag training data.

Jan 2012 — Sep 2012

Peabody, MA

Ph.D. in Computer Engineering, Northeastern University

Dec 2020

- **Dissertation:** It includes 35+ papers with 500+ citations and my involvement in the global community.
- **In summary:** Under [Dr. Yun \(Raymond\) Fu](#), as an integral part of [SMILE Lab](#) and supported as an ALERT DHS STEM Career Development & Research Fellow, researched pattern recognition and data synthesis from various signal types. Cutting-edge work in big data for automatic face understanding (e.g., [1]) by introducing image and video databases with novel deep learning techniques as benchmarks; hosted several data challenge workshops to promote and attract researchers to the problem (e.g., [5]). Besides, research topics included, but were not limited to, big data recognition and understanding, landmark detection, generative modeling, learning with minimal or imbalanced data, bias in ML with consideration to its impact on society, and the construction of multi-modal databases (e.g., [2]). Vast soft skills acquired: being active in diverse groups, team building, handling administration, and measuring practical significance; becoming a competent writer, especially in the form of research and proposals, and an all-around fluent communicator, whether in a technical, formal, and informal, and acquired a vast professional network - familiarized, and often personally, with many experts throughout the technological and entrepreneurial worlds. An experience that helps reveal my passion for learning, teaching, and innovating.
- **Activities:** Ph.D. Council member, ECE Department Representative, student mentor for the Gordon Scholar Program, and the Secretary and later the Research Ambassador Student Research Engagement Committee (SREC), an effort I helped start.
- **Key courses:** Advances in Deep Learning, Applied Probability & Stochastic Processes, Parallel Processing Data Analysis, Advanced Pattern Recognition, Machine Learning, 2D Signal & Image Processing, Human-Centered Computing, Principles of Assistive Robotics, Advanced Computer Vision, Fundamentals of Computer Engineering.

Part-time Faculty, Northeastern University

Jan 2017 – Dec 2020

- Designed curriculum (e.g., syllabus, lectures, assignments) for an undergraduate course, [Comp. Methods for Data Analysis](#).
- Lectured via PowerPoint, board work, code demos, and group activities.
- Emphasized practical use cases and project-based assignments to accelerate the transition from learning to applying.
- Covered vast topics: Python basics, Python for data science (e.g., NumPy, Pandas, Seaborn, SkLearn), exploratory data analysis, data types, regression-based modeling, classification, data visualization, and intro to deep learning.
- Offered supplemental workshops for students throughout the semester: how to build technical power-points, how to write a technical paper, and resume critique (post-semester, as a means to enhance student resumes and experiences).
- Awarded *Best Teacher (2019-20)* of the Department of Electrical and Computer Engineering via the College of Engineering.
- Served as a TA in *Image Processing*, lead lecturer for the *Electronics* lab, and guest lecturer for *Data Visualization*.

Bachelor of Science in Electrical and Computer Engineering, Northeastern University

May 2014

- **In summary:** served as an RA in the [Optical Science Lab](#) and [NUCAR](#), including two summers of [NHS REUs](#). Also, won department-wide [Senior Capstone](#) (ranked 1st) and the [Freshman Remote Control Design Contest](#) (1st place).
- **Key courses:** Robotics, Computer Vision, Computer & Telecomm Networks, Data Visualization, Communication Systems, E-Mags Fields & Waves, Computer Architecture & Organization, Noise & Stochastic Processes, Advanced Tech. Writing, Optimization Methods, Electronics (intro, I, and II), Prof. Issues in Engineering, Algorithms, Data Structure Engineering,

SELECTED PUBLICATION

1. Joseph P Robinson, Ming Shao, and Yun Fu. "Survey on the Analysis and Modeling of Visual Kinship: A Decade in the Making." in *IEEE Trans. on Pattern Analysis and Machine Intelligence* (2021). [[link](#)]
2. Joseph P Robinson, et al. "Families In Wild Multimedia (FIW-MM): A Multi-Modal Database for Recognizing Kinship." in *IEEE Trans. on Multimedia* (2021). [[link](#)]
3. Joseph P Robinson, et al. "Face Recognition: Too Bias, or Not Too Bias?" in *CVPR* (2020). [[link](#)]
4. Joseph P Robinson, Y Li, N Zhang, Y Fu, and S Tulyakov, "Laplace Landmark Localization." in *ICCV* (2019). [[link](#)]
5. Joseph P Robinson, et al. "Visual Kinship Recognition of Families In the Wild" *IEEE TPAMI* (2018). [[link](#)]

For other papers, <https://www.jrobs-vision.com/publications> (or [Scholar](#)); also, check out my blog on [Medium](#).

SKILLS

Machine Learning	Understand, applied, implemented, and analyzed textbook algorithms; invented and published novel methods; comprehensive understanding of deep learning, classic approaches, and statistical modeling; data capture, acquisition, labeling, and evaluation; knowledge of data bias, along with problems of minimal data, and the means to mitigate (e.g., semi-supervision, self-supervision, adversarial learning)
SW Tools and Languages	MATLAB®, Python, GPGPU, Markdown, Pytorch, Tensorflow, Keras, JAVA, CUDA, Spark, Bash, C++, ROS, Gazebo, C, \LaTeX , SolidWorks®, AutoCAD®, Confluence, Notion
HW	Cameras & Video feeds, Depth sensors, Spectrum analyzer, Computer clusters, Function generator, Servos, Photo-detector, Oscilloscope, Multi-meter, RazPi, Board design, Solder, Optical alignment & specs
ML-Ops and DevOps	W&B, Tensorboard, Colab, AWS, Docker, Kubernetes, Colab, Codalab, Jenkins, Git, JIRA
Personal	Communication, Research, Teaching, Mentoring/Advising, Writing, Interpersonal, Leadership, Learning