

^{From} CR/RSI1-NA Our Reference Jonathan Francis Telephone +14123903216 Pittsburgh, PA 13 November 2024

Robot Learning Research Intern – Job Requisition

Description Robot Learning Research Internship

The Robot Learning Lab at Bosch Research Pittsburgh invites knowledgeable research interns for investigations at the intersection of Robotics, Multimodal Machine Learning, Embodied AI, Computer Vision, and Natural Language Processing. We seek to tackle challenging robotics and automation problems, having large-scale industrial impact; we also seek to formulate these industrial problems as interesting and important scientific investigations—often leveraging open-source models, methods, benchmarks, and simulators—with the ultimate goal of deploying these systems to the real world, to augment or work alongside humans and other agents. Multiple members of the lab dual-affiliate with Carnegie Mellon University and, together with collaborators from the Robotics Institute and Language Technologies Institute, we continue to make several key developments in dexterous manipulation, interactive perception for mixed prehensile and non-prehensile manipulation tasks, cross-embodiment transfer learning, few-shot policy generalization through robot trajectory retrieval, unseen / open-vocabulary mobile manipulation, online policy adaptation and failure reasoning through agentic foundation model frameworks, and more.

We expect the intern to display independence and maturity as a researcher, using their experience to construct compelling problem statements, engage in rigorous literature reviews and analyses, design and execute experimental plans, and extract salient insights from the experimental results. To be successful, we expect candidates to have experience in dealing with challenging problems in transfer representation learning and robotics, including: (i) learning safe, robust, or generalizable robot state representations; (ii) designing useful regularization objectives, pretext tasks, or auxiliary objectives; (iii) adapting or transferring representations across different domains (e.g., different embodiments, environments, sim-to-real, tasks, etc.); (iv) dealing with the practicalities related to implementing neural policies, e.g., non-convex optimization "tricks" and multi-machine/multi-GPU parallelized training of large models; (v) conducting careful model performance characterization + error analyses, e.g., determining informative ablations and baselines, inspecting and visualizing learned representations, identifying dataset biases; (vi) using closed- and open-source Vision-Language foundation models, e.g., for perception, planning, world-modeling, progress-monitoring, control, etc.; (vii) fine-tuning foundation models on few-shot examples or large-scale datasets.

Finally, the intern will be expected to contribute to the preparation of industrial patents and to work with teammates to publish a high-quality research paper in a major conference venue.



From	Our Reference	Tel
CR/RSI-NA	Jonathan Francis	+141

Tel +14123903216 Pittsburgh, PA 13 November 2024

Tasks

- Discuss relevant tasks in Robotics / Embodied AI and quickly agree on research topic
- Perform extensive literature review, to understand the current state-of-the-art methods
- Generate an R&D plan—detailing the relevant research questions, selected AI/ML task(s), intended methodology, experiments, evaluation metrics, and publication venue targets
- Maintain an on-going report of progress towards goals and present related literature and project status to supervisor(s) and/or colleagues on a weekly basis
- Design and implement proposed methodology, according to the above research plan, while remaining aware of any new developments in the field
- Perform extensive evaluation of the proposed approach and generate results (e.g., proofs, quantitative results that compare the main results with baselines and ablations, qualitative analysis that visualizations the behavioral tendencies of the approach)
- Work with supervisor(s) to publish a high-quality research paper in a major conference venue, e.g., NeurIPS, ICML, ICLR, CoRL, RSS, CVPR/ICCV/ECCV, ICRA, IROS, etc.

Qualifications

- Strong background in ML, with emphasis on multimodality and/or representation learning
- Strong background in Robot Learning, Robotics, or Embodied AI
- Extensive experience in *from-scratch* neural model implementation, e.g., using PyTorch
- Extensive experience with data analytics toolkits, such as numpy, pandas, and scikit-learn
- Extensive experience in implementing, training/fine-tuning, and evaluating the performance of CV models, NLP models, policies, etc.
- Extensive experience in leveraging Large Language Models, Vision-Language Models, and/or foundation models that are grounded with other modalities (e.g., audio, haptics, etc.)
- Extensive experience in training neural models on multi-machine or multi-GPU setups
- Extensive experience in working on robot hardware platforms
- Extensive publication history in top conference venues; is a mature researcher
- (Preferred) Extensive experience in development in Python on Linux-based systems
- (Preferred) Theoretical background in ML topics, e.g., transfer representation learning, nonconvex optimization, reinforcement learning, provably safe learning, etc.

Other Requirements

- Your degree level: pursuing doctoral degree, or recent post-doctoral researcher
- Your major: Robotics, Machine Learning, Computer Engineering, or related

Logistics

- Internship location: Pittsburgh, Pennsylvania, United States
- *Start date:* The internship typically starts sometime between April and June (flexible)
- Duration: ~12 weeks (extension possible; subsequent research collaboration encouraged)