

## CURRICULUM VITAE

### RESEARCH INTERESTS

I am passionate about enabling intelligent agents, like robots, to adapt and learn continually as they act. I have pursued this goal by investigating how to mitigate forgetting in neural networks exposed to novel data and distribution shift, and how deep reinforcement learning agents can improve their accuracy and learning efficiency in novel environments. I am especially interested in opportunities to enact positive social change in the world leveraging these technologies. My areas of expertise include:

Machine learning (reinforcement learning, imitation learning, continual learning, deep learning, few-shot learning, active learning, learning from demonstration), computer vision, semantic segmentation, natural language processing, human-robot interaction, affordance learning, robotics.

### EDUCATION

*Doctor of Philosophy (in progress)* – **Georgia Institute of Technology**, Atlanta, GA, 2016-Present

- Major: Robotics
  - Core areas: Artificial Intelligence, Perception, Human Robot Interaction
  - Minor: Technology Innovation: Generating Economic Results
- Advisors: Dr. Mark Riedl, Dr. Irfan Essa

*Master of Science* – **University of Pennsylvania**, Philadelphia, PA, 2011-2013

- Major: Robotics
- Advisor: Dr. Kostas Daniilidis

*Bachelor of Science* – **Georgetown University**, Washington, D.C., 2007-2011

- Major: Physics, Mathematics

### REFEREED PUBLICATIONS

1. **Balloch, J. C.**, Lin, Z., Hussain, M., Srinivas, A., Wright, R., Peng, X., Kim, J., Riedl, M. "NovGrid: A Flexible Grid World for Evaluating Agent Response to Novelty." *AAAI2022 Spring Symposium on Designing Artificial Intelligence for Open Worlds*. (2022) **[Long Oral]**
2. Smith, J., **Balloch, J. C.**, Hsu, Y. C., Kira, Z. "Memory-efficient semi-supervised continual learning: The world is its own replay buffer." In *2021 International Joint Conference on Neural Networks (IJCNN)*, IEEE, (2021).
3. Castricato, L., Frazier, S., **Balloch, J. C.**, Riedl, M. "Fabula Entropy Indexing: Objective Measures of Story Coherence." In *Proceedings of the Third Workshop on Narrative Understanding*, pp. 84-94. (2021).
4. Castricato, L., Frazier, S., **Balloch, J. C.**, Riedl, M. "Tell Me A Story Like I'm Five: Story Generation via Question Answering". *Proceedings of the 3rd Workshop on Narrative Understanding* (2021).

5. Smith, J., Hsu, Y. C., **Balloch, J. C.**, Shen, Y., Jin, H., and Kira, Z. "Always be dreaming: A new approach for data-free class-incremental learning." In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9374-9384. (2021).
6. Peng, X., **Balloch, J. C.**, Riedl M. "Detecting and Adapting to Novelty in Games." *AAAI2020 Workshop on Reinforcement Learning in Games*. (2021).
7. Banerjee, S., Daruna, A., Kent, D., Liu, W., **Balloch, J. C.**, Jain, A., Krishnan, A., Chernova, S. "Taking Recoveries to Task: Recovery-Driven Development for Recipe-based Robot Tasks." *IEEE International Symposium on Robotics Research*, (2019).
8. Nair, Lakshmi, **Balloch, J. C.**, Chernova, S. "The MacGyverbot: Tool Construction by Autonomous Agents." *IEEE International Conference on Robotics and Automation*, 2019.
9. **Balloch, J. C.**, Chernova, S. "An RGBD segmentation model for robot vision learned from synthetic data." *Robotics Science and Systems (RSS): Workshop on Spatial-Semantic Representations in Robotics*, (2017).
10. Endo, Y., **Balloch, J.**, Grushin, A., Lee, M.W., Handelman, D. "Landmark-Based Robust Navigation for Tactical UGV Control in GPS-Denied Communication-Degraded Environments." *SPIE Unmanned Systems Technology XVIII*, (2016).
11. West, R. A., Ovanessian, A., Turtle, E. P., Ray, T., **Balloch, J.**, Dumont, P., Lavvas, P., Lorenz, R., Rannou, P. "Titan's Detached Haze and Polar Vortex: Large-Amplitude Seasonal Variations." *Lunar and Planetary Science Conference*, 43, (2012).
12. West, R. A., **Balloch, J.**, Dumont, P., Lavvas, P., Lorenz, R., Rannou, P., Turtle, E. P., Ray, T. "The Evolution of Titan's detached haze layer near equinox in 2009." *Geophysical Research Letters*, 38, doi: 10.1029/2011GL046843, (2011).

## OTHER PUBLICATIONS

1. Castricato, L., Frazier, S., **Balloch, J. C.**, Tarakad, N., Riedl, M. "Automated Story Generation as Question-Answering." *arXiv preprint arXiv:2112.03808* (2021).
2. Peng, X., **Balloch, J. C.**, Riedl M. "Detecting and Adapting to Novelty in Games." *ICSR2020 Workshop on Creativity*. (2020).
3. **Balloch, J. C.**, Aggraval, V., Essa, I., Chernova, S. "Unbiasing Semantic Segmentation for Robot Perception using Synthetic Data Feature Transfer." [arXiv:1809.03676](https://arxiv.org/abs/1809.03676), (2018).

## AWARDS

- Public Interest Technology Universities Network (PITUN) Fellowship**, 2020-2021
- ◆ Two-semester fellowship program partnering with social scientists at Georgia State University professor William Sabol. Focused on interdisciplinary projects addressing historic inequity challenges.
  - ◆ We compare the effectiveness of remote community supervision meetings in Georgia using deep learning-based multi-modal sentiment analysis. We examine the relationship between sentiment, race, and recidivism.
  - ◆ Support from Ford Foundation, New America, and the Hewlett Foundation.
- TI:GER Fellowship (GT Scheller School of Business)**, 2020-Present

- ◆ Georgia Tech MBA and PhD students from engineering, computing, and science colleges work to translate ideas into successful technology innovations.
- ◆ Educational program for the [Creative Destruction Lab](#) (CDL-Atlanta), a global milestone-based mentoring program for scalable science-based startups.

## PROFESSIONAL AND RESEARCH EXPERIENCE

**Graduate Researcher – [Georgia Institute of Technology](#)** 2016-Present

- ◆ *Current research project:* Enhancing guided reinforcement learning using natural language explanations.
  - Prior work has shown that agents learn more efficiently by reinforcement when guided by the advice of experts. We explore whether the explanations of the expert actions can be used alongside the advice. We propose that a recurrent 'rationale module' that learns a transformation between the state and action of a policy to the same semantic space as a natural language action explanation.
  - Exploring extensions to this work using explanations to improve the generalizability of policies learned from imitation learning. Tested in the MALMO Minecraft environment.
- ◆ *Current research project:* Mitigating catastrophic forgetting with semi-supervised continual learning using locally occurring unsupervised data.
  - Neural networks models exhibit catastrophic forgetting when data shifts during learning, as in the case when new classes are introduced after the model has been trained. Most people attempt to correct this by memorizing old samples, but this is not scalable.
  - Combining concepts from out of distribution detection, continual learning, and semi-supervised learning, we develop a memory-free solution that performs comparably to methods that memorize samples from prior tasks. Tested on MNIST, CIFAR100, and Matterport.
- ◆ *Prior project:* The MacGyverbot: tool construction by autonomous agents.
  - Given a robotics task, a reference tool that is not present, and random objects in the environment, our approach enables a robot to construct a new tool that can also complete the task by imitating the geometry of the reference tool.
  - Paper accepted for publication and presentation at ICRA 2019.
- ◆ *Prior project:* Unbiasing semantic segmentation for real-time robot perception using synthetic data pretraining.
  - Showed that pretraining on a large amount of synthetic data improves performance on real data compared to comparable amounts of real pretraining data by reducing the bias toward the pretraining dataset.
  - Work presented at workshop at 2017 Conference on Robotics: Science and Systems, full paper submitted to ArXiv in 2018.

## PhD Intern – SRI International

2021

- ◆ Worked with the Autonomous Agents team at in the AI group of Palo Alto.
- ◆ Research project: Generating and learning behavior trees using neural fictitious self-play for interpretable strategy in multi-agent reinforcement learning.
  - Technique increased convergence efficiency to winning policy on LaserTag and internal Battlefield simulators.
  - Still in use by SRI for grants with the military for officer training in Red Team Doctrine exercises.

## Graduate SWE Intern – Google

2018

- ◆ Worked with the Mobile Vision Research team at Google AI in Seattle.
- ◆ Research project: actively construct mini-batches online for more efficient training of machine learning models optimized with stochastic gradient descent (SGD).
  - Preliminary results showed a 2x reduction in training time to convergence on classification tasks with MNIST and CIFAR10.

## Robotics Engineer – Intelligent Automation, Inc.

2013-2016

- ◆ Specialized in design and development of computer vision, sensor fusion, and control systems on DARPA and DoD robotics research and development projects, collaborating with both industry and academic groups.
- ◆ Led computer vision research and development for Bearing-based Landmark Navigation robotic system.
  - Improved contour-based feature tracking persistence from <75% to 94% and optimized code to increase run-time efficiency from 0.4 fps to 5 fps with HD streaming input.
  - Designed omnidirectional camera sensor head from four COTS cameras with IMU and a 3D printed mount, and developed API for integration with our robot platform.
- ◆ Led an effort working with a dual-arm manipulator robot for learning autonomous furniture assembly of IKEA furniture from demonstration.
- ◆ Implemented ROS (Robot Operating System) integration as part of the AEODRS protocol for a multi-arm mobile robot system.
  - Demonstrated ease-of-use for multi-arm control at the DARPA Robotic Challenge Trials Expo.

## Graduate Student Researcher – GRASP Lab, University of Pennsylvania

2012-2013

- ◆ Worked with the TROOPER team, a collaboration between Lockheed Martin and UPenn, on the DARPA Robotics Challenge.
- ◆ Enabled bipedal walking and standing stability in simulation by implementing impedance control and ZMP algorithms in ROS using C++.

# Jonathan C. Balloch

Webpage: [jballoch.com](http://jballoch.com)  
E-mail: [jon.balloch@gmail.com](mailto:jon.balloch@gmail.com)  
Mobile: (203) 912-1793

---

## Graduate Research Intern – Lockheed Martin, 2012

- ◆ Worked with the Phenomenology team at the Advanced Technology Center in Palo Alto.
- ◆ Developed a MATLAB package which reduced digital noise and increased accuracy in laser simulations.

## Planetary Science Intern – NASA Jet Propulsion Laboratory, 2010

- ◆ Funded through the NASA Space Grant.
- ◆ Modeled radiative transfer in Titan's detached haze layer in FORTRAN.
- ◆ Discovered the cyclic seasonal collapse and expansion in altitude and eccentricity of Titan's atmosphere.
- ◆ Published findings in *Geophysical Research Letters*.

## TECHNICAL SKILLS

Python, SciPy Stack, PyTorch, TensorFlow, LaTeX, ROS, Gazebo, Linux shell, C++, Caffe, OpenCV, Boost, C, MATLAB, Arduino, PCL, Blender, Java, Qt, Android, SolidWorks, C, Mathematica

## TEACHING EXPERIENCE

CS6476: Computer Vision – <u>Georgia Institute of Technology (OMSCS)</u>	2020-2022
CS7641: Machine Learning – <u>Georgia Institute of Technology (OMSCS)</u>	2020
CS3600: Artificial Intelligence – <u>Georgia Institute of Technology</u>	2020
CS4476: Computer Vision – <u>Georgia Institute of Technology</u>	2019
CS7642: Reinforcement Learning – <u>Georgia Institute of Technology (OMSCS)</u>	2019
CS6476: Computer Vision – <u>Georgia Institute of Technology</u>	2017
CS4641: Machine Learning – <u>Georgia Institute of Technology</u>	2016
ESE-505: Introduction to Control Systems – <u>University of Pennsylvania</u>	2013
MEAM-510: Mechatronic Systems – <u>University of Pennsylvania</u>	2012
PHYS-252: Electricity & Magnetism (Lecturer) – <u>Georgetown University</u>	2010-2011
PHYS-101/102: Intro to Physics – <u>Georgetown University</u>	2009-2011

## VOLUNTEER LEADERSHIP EXPERIENCE

Vice President, RoboGrads Student Organization – <u>Georgia Institute of Technology</u> ,	2017-2019
FIRST Robotics Mentor – <u>Team 449</u> , <u>Montgomery Blair High School</u> ,	2014-2016