

Liam Dugan

ldugan@seas.upenn.edu | [linkedin.com/in/liam-dugan](https://www.linkedin.com/in/liam-dugan) | liamdugan.com

EDUCATION

UNIVERSITY OF PENNSYLVANIA	Philadelphia, PA
Ph.D, Computer Science	Aug. 2021 - Present
<i>Advisor: Dr. Chris Callison-Burch</i>	
M.S.E, Robotics (GPA: 3.80/4.00)	Aug. 2017 – Dec. 2020
B.S.E, Computer Engineering & East Asian Studies (GPA: 3.63/4.00)	Aug. 2015 – Aug. 2020
DOSHISHA UNIVERSITY 同志社大学	Kyoto, Japan
(Study Abroad) Kyoto Consortium of Japanese Studies (GPA: 3.70/4.00)	Jun. 2017 – Aug. 2017

RESEARCH EXPERIENCE

Roblox Research	San Mateo, CA
<i>PhD Research Intern</i>	Jun. 2022 (Incoming)
<ul style="list-style-type: none">Expected to work with research team developing novel speech to speech translation systems	
John's Hopkins University	Baltimore, MD
<i>Visiting Research Scholar</i>	Jun. 2021 - Aug. 2021
<ul style="list-style-type: none">Was part of the Machine Translation team at HLTCOE working on customizing MT for Cross-Lingual IR.Developed a novel way of fine-tuning MT models for IR by ignoring non-content words in training dataImprovement in NDCG' observed in Russian, Farsi, and Chinese to English pairs. Publication forthcoming	
University of Pennsylvania	Philadelphia, PA
<i>Graduate Research Assistant</i>	Jun. 2019 – Aug. 2021
<ul style="list-style-type: none">Led the Real or Fake Text project http://roft.io which measures how well humans can detect generated text.Extensive experience with natural language generation systems such as GPT-2, GPT-3, CTRL, GROVER, and T5Used BERT classifiers and Adversarial Decomposition to predict English formality by training on JapaneseWorked on the BETTER project using Machine Translation for Arabic-English cross-lingual event extraction	

WORK EXPERIENCE

NVIDIA Corporation	Santa Clara, CA
<i>Autonomous Driving Software Intern</i>	Jun. 2019 – Aug. 2019
<ul style="list-style-type: none">Assisted development of a platform for on-demand downloads of self-driving car softwareCustom docker images are requested by engineers through a Jenkins server and images come pre-installed with latest software and can then be flashed onto vehicle hardwareService was deployed to over 500 developers on the NVIDIA DRIVE Team	
Robotic Research LLC	Clarksburg, MD
<i>Software Engineering Intern</i>	Jun. 2018 – Aug. 2018
<ul style="list-style-type: none">Worked on Velodyne VLP-16 LIDAR at the driver level for Autonomous Ground Resupply convoysDeveloped and prototyped novel object classifiers for sun speckles, dust, and vegetationUsed a PCA-based volumetric analysis to tag neighboring points in a point cloud as possible vegetation or humans	

TECHNICAL SKILLS

Natural Languages: English (native), Japanese (advanced - 5+ years), Korean (elementary - 0.5 years)
Programming Languages: Python, C/C++, bash, CUDA, Java, JavaScript, Go, HTML/CSS, Verilog, MATLAB
Frameworks: PyTorch, Tensorflow, OpenCV, DXR, Vulkan, OpenGL, WebGL, React, Node, Gatsby, Django
Developer Tools: Git, Docker, Google Cloud Platform, VS Code, emacs, Atom, tmux

PUBLICATIONS

Daphne Ippolito*, **Liam Dugan***, Arun Kirubarajan, Sherry Shi, Chris Callison-Burch. Investigating Human Ability to Detect Boundaries Between Human-Written and Machine-Generated Text. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. (In Preparation), 2022

Daphne Ippolito, **Liam Dugan**, Emily Reif, Ann Yuan, Andy Coenen, Chris Callison-Burch. The Case for a Single Model that can Both Generate Continuations and Fill in the Blank. In *Proceedings of the 2022 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. (To Appear), 2022

Liam Dugan, Eleni Miltsakaki, Shriyash Upadhyay, Etan Ginsberg, Hannah Gonzalez, DaHyeon Choi, Chuning Yuan, and Chris Callison-Burch. A feasibility study of answer-unaware question generation for education. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1919–1926, Dublin, Ireland, May 2022. Association for Computational Linguistics

Liam Dugan. Learning Formality from Japanese-English Parallel Corpora. Master’s thesis, University of Pennsylvania, December 2020

Liam Dugan*, Daphne Ippolito*, Arun Kirubarajan*, and Chris Callison-Burch. RoFT: A Tool for Evaluating Human Detection of Machine-Generated Text. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 189–196, Online, October 2020. Association for Computational Linguistics

Zhengyi Luo, Austin Small, **Liam Dugan** and Stephen Lane. Cloud Chaser: Real Time Deep Learning Computer Vision on Low Computing Power Devices. *Eleventh International Conference on Machine Vision (ICMV 2018)*, Mar 2019

TEACHING EXPERIENCE

Teaching Assistant

- CIS700 (Interactive Fiction & Text Generation) – Spring 2022
- CIS565 (GPU Programming & Architecture) – Fall 2021
- CIS530 (Computational Linguistics) – Spring 2020, Fall 2020
- CIS380 (Operating Systems) – Fall 2018, Spring 2019, Fall 2019 (Head TA)
- CIS240 (Introduction to Computer Systems) – Spring 2017, Fall 2018, Spring 2018

Lectures

- CIS565: “Optimizing Machine Learning with CUDA” – Fall 2021
- CIS565: “Introduction to Machine Learning” – Fall 2021
- CIS380: “Linux Page Replacement Algorithms and Belady’s Anomaly” – Fall 2019

Authored Homework Assignments

- CIS530: “HW7: Transformers and State-of-the-Art Language Models” – Nov. 2020
- CIS530: “HW10: Neural Machine Translation” (with Li “Harry” Zhang) – Apr. 2020

FELLOWSHIPS, AWARDS AND HONORS

Fellowships

- (October 2018) FLAS: Foreign Language and Area Studies Undergraduate Fellowship (East Asia) – \$15,000

Certifications

- (May 2020) Penn Certificate of Japanese Proficiency
- (January 2018) Japanese Language Proficiency Test: N2

Academic Honors

- (2016-2017; 2017-2018) Dean’s List x2

Academic Awards

- (May 2020) Penn Engineering Exceptional Service Award
- (May 2020) Moore School Council Cwikla Award (Most Improved Student) [Nominated]
- (May 2019) Penn Engineering Computer Science Senior Design Award Third Prize

Hackathon Awards

- (February 2018) Most Innovative Use of Technology: Wharton Undergrad FinTech (WUFT) Hacks
- (January 2018) Grand Prize: PennApps XVII
- (January 2018) Best use of Cloud Hosting: PennApps XVII
- (September 2017) Third Prize: PennApps XVI

Invited Talks

- “Intro to Machine Learning and AI Research” - St. Joe’s Prep High School, Philadelphia PA, Feb. 2022
- “Can Humans Detect Boundaries Between Human-Written and Machine-Generated Text?” - University of Pennsylvania, Philadelphia PA, Jan. 2022

Posters

- “A Feasibility Study of Answer-Agnostic Question Generation for Education” - Poster Presentation ACL 2022
- “RoFT: A Tool for Evaluating Human Detection of Machine Generated Text” - Poster Presentation EMNLP 2020

Presentations

- “Learning Formality from Japanese-English Parallel Corpora” - Live Presentation, Master’s Thesis ([video](#))
- “Scene++ VR” - Live Presentation, Penn Engineering Senior Design Demo Day ([video](#))
- “Cloud Chaser” - Live Presentation, PennApps XVII Closing Ceremony ([video](#))
- “Todd: The Inter-Dimensional Robot” - Live Presentation, PennApps XVI Closing Ceremony ([video](#))

PROJECTS

SCENE++ VR | *Oculus Rift, ZED Mini, Unity 3D, Python, YOLOv3, Paperspace* May 2019

- **Won 3rd Prize in Penn Computer Science Senior Design Competition**
- We developed a Unity API that allows VR & AR Developers to query real-world objects around the user
- Hardware: Oculus Rift headset with head-mounted ZED mini depth camera for pass-through Augmented Reality
- We use Spatial Feature Mapping of environment to allow localization and stabilization of queried objects in depth
- Offloading object detection to cloud server allows Scene++ to run on any platform with *virtually no drop in FPS*

RTX EXPLORE | *C++, DirectX Raytracing, NVIDIA TitanV* December 2018

- (**40+ stars on GitHub**) Built the first open-source path tracer in the DirectX Raytracing GPU framework
- Features include: Dynamic model loading from .gltf and .obj, support for texture and normal maps, live editing of scene transformations through GUI interface, specular/refractive/dispersive/transmissive materials, subsurface scattering, anti-aliasing, depth of field

BANKING WITH A VISION | *Python, Javascript, TCP, Bootstrap* February 2018

- **Won Most Innovative use of Technology at WUFT Hacks**
- Use facial key-point mapping algorithm to perform face recognition through front facing webcam
- Faces are used to access database of customer information to save bank tellers having to pull up user information

CLOUD CHASER | *Python, C, TCP, AWS, YOLOv3, Alexa, Raspberry Pi, 3D Printing* January 2018

- **Won Grand Prize and Best use of Cloud Hosting at PennApps XVII**
- Presented a platform that allows low resource IoT devices to do high level image processing on the cloud
- Hardware: Raspberry Pi + camera, 3D printed robot chassis & camera mount, 4 servo motors, Amazon Echo Dot
- Built robot “Chase” to demonstrate our platform. Commands are given to Chase through Echo Dot
- Paper outlining our techniques to reduce latency of streaming for IoT image detection accepted to ICMV 2018

TODD: THE INTER-DIMENSIONAL ROBOT | *C, HC-05 Bluetooth, Arduino, Unity 3D* September 2017

- **Won Third Prize at PennApps XVI**
- Made multiplayer game where player controlling Todd has to dodge objects only visible in virtual world
- Hardware: Arduino, breadboard, 2 servo motors, Bluetooth HC-05 controller
- USB connected Bluetooth controller allows Unity to communicate with “Todd” the robot

REFERENCES

Chris Callison-Burch, Associate Professor
Natural Language Processing
ccb@cis.upenn.edu

Boon Thau Loo, RCA Professor of Computer Science
Director - UPenn Distributed Systems Lab
boonloo@seas.upenn.edu