

Daphne Ippolito

daphnei@seas.upenn.edu

Education

2015–present	PhD in Computer Science expected Computer Science, School of Engineering and Applied Science University of Pennsylvania <i>Expected graduation Spring 2022</i>
2010–2015	BSc, Specialist in Computer Science Minor in Geographic Information Systems University of Toronto

Thesis Research

Advisors: Dr. Chris Callison-Burch at University of Pennsylvania and Dr. Douglas Eck at Google Research

I research the tradeoffs and limitations of generating text with neural language models, as well as strategies for evaluating natural language generation systems. I also study methods for incorporating AI-in-the-loop language generation into assistive tools for writers.

Publications in Conference Proceedings and Journals

- 2021 Ann Yuan, Daphne Ippolito, Vitaly Nikolaev, Chris Callison-Burch, Andy Coenen, and Sebastian Gehrmann. “SynthBio: A Case Study in Curation of Text Datasets.” NeurIPS Datasets and Benchmarks track. December 2021.
- 2021 Yoshua Bengio, Daphne Ippolito, Richard Janda, Max Jarvie, Benjamin Prud’homme, Jean-Francois Rousseau, Abhinav Sharma, Yun William Yu, “Inherent privacy limitations of decentralized contact tracing apps.” *Journal of the American Medical Informatics Association*. January 2021.
- 2020 Liam Dugan*, Daphne Ippolito*, Arun Kirubakaran*, Chris Callison-Burch “RoFT: A Tool for Evaluating Human Detection of Machine-Generated Text.” EMNLP 2020 Demo track.
- 2020 Yoshua Bengio, Richard Janda, Yun William Yu, Daphne Ippolito, Max Jarvie, Dan Pilat, Brooke Struck, Sekoul Krastev, Abhinav Sharma, “The need for privacy with public digital contact tracing during the COVID-19 pandemic.” *The Lancet Digital Health*. July 2020.
- 2020 Daphne Ippolito, David Graingier, Douglas Eck, Chris Callison-Burch. “Toward Better Storylines with Sentence-Level Language Models” ACL 2020.
- 2020 Daphne Ippolito*, Daniel Duckworth*, Chris Callison-Burch, Douglas Eck, “Automatic Detection of Generated Text is Easiest when Humans are Fooled.” ACL 2020.
- 2019 Daphne Ippolito*, Reno Kriz*, Maria Kustikova, Joao Sedoc, Chris Callison-Burch. “Comparison of Diverse Decoding Methods from Conditional Language Models.” ACL 2019.
- 2019 Joao Sedoc*, Daphne Ippolito*, Arun Kirubakaran, Jai Thirani, Lyle Ungar, Chris Callison-Burch. “Chateval: A tool for chatbot evaluation.” NAACL 2019 Demo track.
- 2018 Andrew Jaegle, Stephen Phillips, Daphne Ippolito, Kostas Daniilidis. “Understanding Image Motion with Group Representations.” ICLR 2018.

- 2018 | John Hewitt*, Daphne Ippolito*, Brendan Callahan, Reno Kriz, Derry Wijaya, Chris Callison-Burch. “Learning Translations via Images with a Massively Multilingual Image Dataset.” ACL 2018.

*Co-first authors

Publications in Refereed Workshops

- 2021 | Andy Coenen*, Luke Davis*, Daphne Ippolito*, Emily Reif*, Ann Yuan*, “Wordcraft: a Human-AI Collaborative Editor for Story Writing.” HCI+NLP Workshop at EACL 2021.
- 2019 | Daphne Ippolito, David Grangier, Chris Callison-Burch, Douglas Eck, “Unsupervised Hierarchical Story Infilling.” First Workshop on Narrative Understanding at NAACL 2019.
- 2018 | Daphne Ippolito, Anna Huang, Curtis Hawthorne, Douglas Eck, “Infilling piano performances.” Workshop on Machine Learning for Creativity and Design at NeurIPS 2018.
- 2018 | Anna Huang, Curtis Hawthorne, Daphne Ippolito, Douglas Eck, “Transformer-NADE for Piano Performances.” Workshop on Machine Learning for Creativity and Design at NeurIPS 2018.

*Authors listed alphabetically

Other Articles and Preprints

- 2021 | Katherine Lee†, Daphne Ippolito†, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. “Deduplicating Training Data Makes Language Models Better” *In Submission*.
- 2021 | Emily Reif†, Daphne Ippolito†, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei. “A Recipe For Arbitrary Text Style Transfer with Large Language Models” *In Submission*.
- 2021 | Daphne Ippolito, Liam Dugan, Emily Reif, Ann Yuan, Andy Coenen, and Chris Callison-Burch. “Continuation is a Sub-Task of Fill in the Blank: Why Not Train for Both?” *In Submission*.
- 2020 | Hannah Alsdurf*, Yoshua Bengio*, Tristan Deleu*, Prateek Gupta*, Daphne Ippolito*, et al. “COVI White Paper.” arXiv preprint:2005.08502.
- 2020 | Hyunghoon Cho, Daphne Ippolito, Yun William Yu, “Contact Tracing Mobile Apps for COVID-19: Privacy Considerations and Related Trade-offs .” arXiv preprint:2003.11511.
- 2020 | Hugh Zhang, Daniel Duckworth, Daphne Ippolito, Arvind Neelakantan, “Trading Off Diversity and Quality in Natural Language Generation.” arXiv preprint:2004.10450.

† Co-first authors

* Authors listed alphabetically

Teaching Experience

Spring 2020	Co-Instructor for CIS700, <i>Interactive Fiction and Text Generation</i> With Chris Callison-Burch, University of Pennsylvania
Spring 2018	Teaching Assistant for CIS530, <i>Computational Linguistics</i> Taught by Chris Callison-Burch, University of Pennsylvania
Spring 2017	Teaching Assistant for CIS580, <i>Machine Perception</i> Taught by Jianbo Shi, University of Pennsylvania
Fall 2016	Teaching Assistant for CIS390, <i>Robotics: Planning and Perception</i> Taught by Kostas Daniilidis, University of Pennsylvania
Spring 2015	Teaching Assistant for CSC209, <i>Software Tools and Systems Programming</i> Department of Computer Science, University of Toronto
Fall 2014	Teaching Assistant for CSC104, <i>Computational Thinking</i> Department of Computer Science, University of Toronto

Industry and Other Experience

2021-present	Research Scientist in Google Brain I currently work part-time as a research scientist on the Language team at Google Brain.
2018-2021	Student Researcher in Google Brain My PhD research was sponsored by a student researcher position at Brain where I have worked closely with Magenta, a team focused on music generation, and with the Brain Language team.
Summer 2017	Google Software Engineering Internship in Research I worked with the Machine Perception team to use self-supervised learning techniques for which there is nearly unlimited data to help with supervised tasks in which dataset sizes are limited.
Summer 2015	Google Software Engineering Internship I worked on the Google Voice Actions team, making improvements to an API used by third-party developers to design custom voice actions.
August 2016	Research Assistant at Wonderwerk Cave Advisor: Dr. Michael Chazan, Department of Anthropology, University of Toronto I participated in archaeological research at a prehistoric cave site in South Africa with some of the oldest evidence for human-controlled fire. I set up a protocol for photogrammetry, and used a drone to map illegal diamond mining at a heritage site.
2013-2014	Software Developer at WaveDNA Professional Experience Year, 12 month internship I worked on Liquid Rhythm, drum beat generation software that offers improved ways for electronic musicians to compose MIDI drum beats. I gained expertise in Java application development and programmatically interacting with MIDI and audio.
May-Dec 2012	NSERC-Create Award in Auditory Cognitive Neuroscience Advisor: Dr. Laurel Trainor, Department of Psychology, McMaster University Canada's version of NSF REU. I built software and hardware for studying infant perception, allowing researchers to design experiments examining reactions to auditory and visual stimuli.

Honors and Awards

October 2021	EECS Rising Stars Workshop I was selected to participate in MIT's Rising Stars academic workshop for graduating students from underrepresented groups interested in pursuing academic careers in computer science.
2016	NSF GRFP Honorable Mention I received an honorable mention for the NSF Graduate Research Fellowship Program.
2015-2016	Integrative Graduate Education and Research Traineeship An NSF traineeship awarded to PhD students pursuing cross-disciplinary research spanning engineering and the social sciences.
Spring 2015	Winner of the Level Up Game Showcase I worked on a team of six to develop Pitfall Planet, a 3D cooperative puzzle game set in space. Pitfall Planet won 1st place out of over 100 game entries at the Level Up Showcase.
Fall 2014	Konrad Group Award by University of Toronto A prize awarded yearly to recognize a student who demonstrates creativity and innovation with respect to the field of technology.
January 2014	Great Canadian Appathon Sponsored by XMG Studios I worked on a team of four to develop a mobile game in 48 hours. Our game was ranked in the top 15 out of 140 submissions across Canada, and won the award for Best Art & Aesthetics.
August 2012	Scholarship Recipient for Independent Experiential Study Program Advisor: Dr. Edward Swenson, Department of Anthropology, University of Toronto I was funded to participate in the excavation of a Moche ceremonial centre in northern Peru. I then wrote ArcheoPhoto, an app facilitating metadata annotations on excavation photos.
2011	American Schools of Oriental Research, Platt Excavation Fellowship This competitive award supported my participation in an archaeological field school at the site of Huqoq in northern Israel.

Invited Talks

2021	“Language Models Memorize their Training Data” Computational Linguistics Lunch (CLUNCH) at University of Pennsylvania
2020	“Controllable Text Generation with Neural Networks and its Limitations” Invited talk to the Department of Social Sciences and Humanities at ETH Zurich
2020	“The Pitfalls of Selecting a Decoding Strategy” Guest Lecture in Johns Hopkins' EN.601.767, <i>Deep Learning Methods for Automated Discourse</i>
2019	“Detecting Machine-Generated Text” Computational Linguistics Lunch (CLUNCH) at University of Pennsylvania
2018	“Clustering Word Senses using Image Features” Mid-Atlantic Student Colloquium on Speech, Language and Learning (MASC-SLL)

Service

July 2022	NAACL Student Research Workshop I am one of the lead organizers for the Student Research Workshop at NAACL 2022.
December 2021	Workshop on Machine Learning for Creativity and Design I am a co-organizer for the 2021 iteration of the NeurIPS Creativity Workshop.
May 2021	Workshop on Enormous Language Models I co-organized the WELM workshop at ICLR 2021. I participated in designing the workshop's format (WELM was a new workshop with no precedents), inviting and organizing invited speakers, and managing the Zoom meeting.
2020-2021	School of Engineering and Applied Science Climate Committee I was one of two graduate student representatives to a committee working to improve the cultural climate at Penn Engineering.
December 2020	Workshop on Machine Learning for Creativity and Design I co-organized the Creativity workshop at NeurIPS 2020. I participated in the selection of invited speakers and accepted papers, and I ran the Discord server for the virtual event.
2019-2021	Dean's Doctoral Advisory Board I was a member of the student board which advises the Dean of Engineering on the experiences and needs of PhD students across the School of Engineering.
2017-present	CIS Doctoral Association I co-founded the student government association for PhD students in my department and have been a member of the leadership team since then.
May 2017	Mid-Atlantic Computer Vision Workshop I organized the 2017 MACV workshop at UPenn. I coordinated speakers and workshop logistics, and acted as main point of contact for attendees.

Other Activities

2015–present	University of Pennsylvania Music Ensembles I have played flute and piccolo in Penn's top audition-only Symphony Orchestra and bass clarinet in the Wind Ensemble. I currently play flute in the Baroque Ensemble.
2015–2018	FIRST Robotics Mentor and Robot Design Judge I taught programming and engineering skills to a novice FIRST Robotics team at an under-served primary school in Philadelphia, and I regularly volunteered as a competition judge.
October 2013	OurCS – Workshop for Undergraduate Women in Computer Science School of Computer Science, Carnegie Mellon I was selected to participate in a four-day research workshop at Carnegie Mellon aimed to improve gender balance in computing research.
2011–2015	Hart House Orchestra I played flute and piccolo in a large audition-only community orchestra at the University of Toronto.