

Jamie Hayes

CONTACT INFORMATION

Email: j.hayes@cs.ucl.ac.uk

WWW: <https://jhayes14.github.io>

EXPERIENCE

Google DeepMind, London, UK

Research Scientist

May, 2020 -

Research into verification and robustness guarantees in machine learning.

Google DeepMind, London, UK

Internship

May, 2019 - September, 2019

Research into zero-bit watermarking applications in machine learning.

Google Research, Mountain View, CA, USA

Internship

July, 2018 - September, 2018

Research into generative and adversarial machine learning. Developed new techniques for unsupervised style transfer, and new probabilistic conditioning methods for generative models.

Microsoft Research, Cambridge, UK

Internship

February, 2018 - May, 2018

Research into security and privacy attacks and defenses in multi-party machine learning.

Naval Research Laboratory, Washington, DC, USA

Internship

August, 2017 - October, 2017

Research into adversarial machine learning and network traffic analysis. Proposed a project on using machine learning techniques for performing privacy-preserving experiments into live traffic analysis attacks on Tor.

Government Digital Services, London, UK

Intern for the SecOps team

March, 2017 - July, 2017

I developed and implemented a privacy-preserving machine learning pipeline to aid threat analysis and improve Transaction Monitoring (TxM) on the GOV.UK Verify system.

University of Manchester, Manchester, UK

Network Technician

September, 2011 - March, 2013

Maintenance of the internal University network that provided a connection to over 10,000 students. Duties included server maintenance, switch configurations, some SDN programming.

EDUCATION

Dept. of Computer Science, University College London, UK

Sep. 2014 - Present - Ph.D. candidate

Privacy, Security and Machine Learning

Advisor: Prof. George Danezis

Second Advisor: Prof. Thore Graepel

Computer Laboratory, University of Cambridge, UK

Sep. 2013 - Apr. 2014 - Researcher

Parameterized Computational Complexity of the Graph Isomorphism Problem

Advisor: Prof. Anuj Dawar

Dept. of Mathematics, University of Manchester, UK

Sep. 2007 - Jun. 2011 - Master of Mathematics

First Class Grade (graduated 4th out of a class of 76)

Advisor: Prof. Richard Sharp

HONORS AND
AWARDS

NeurIPS 2018 Student Travel Award

NeurIPS 2017 Student Travel Award

Google Phd Fellowship in Machine Learning (2017-2020)

Invited to Google PhD Summit in Security (2016) and Machine Learning (2017)

Academic Center of Excellence Studentship, 2014

Engineering and Physical Sciences Research Council (EPSRC) Doctoral Training Studentship, 2013

RESEARCH

Extensions and limitations of randomized smoothing for robustness guarantees

J Hayes 06-2020 CVPR (workshop track)

A framework for robustness certification of smoothed classifiers using f-divergences

K Dvijotham, J Hayes, B Balle, Z Kolter, C Qin, A Gyorgy, K Xiao, S Goyal, P Kohli 05-2020 ICLR

Towards transformation-resilient provenance detection of digital media

J Hayes, K Dvijotham, Y Chen, S Dieleman, P Kohli, N Casagrande 09-2019

LOGAN: Membership inference attacks against generative models

J Hayes, L Melis, G Danezis, E De Cristofaro 07-2019 PETS

A note on hyperparameters in black-box adversarial examples

J Hayes 12-2018

Evading classifiers in discrete domains with provable optimality guarantees

B Kulynych, J Hayes, N Samarin, C Troncoso 12-2018 NeurIPS (workshop track)

Contamination attacks in multi-party machine learning

J Hayes, O Ohrimenko 12-2018 NeurIPS

On visible adversarial perturbations & digital watermarking

J Hayes 06-2018 CVPR (workshop track)

Learning universal adversarial perturbations with generative models

J Hayes 05-2018 DLS (IEEE S&P workshop)

Generating steganographic images via adversarial training

J Hayes, G Danezis 12-2017 NeurIPS

AnNotify: A private notification service
A Piotrowska, J Hayes, N Gelernter, G Danezis, A Herzberg 10-2017 WPES

The Loopix anonymity system
A Piotrowska, J Hayes, T Elahi, S Meiser, G Danezis 08-2017 USENIX Security

Website fingerprinting defenses at the application layer
G Cherubin, J Hayes, M Juarez 07-2017 PETS

TASP: Towards anonymity sets that persist
J Hayes, C Troncoso, G Danezis 10-2016 WPES

k -fingerprinting: a robust scalable website fingerprinting technique
J Hayes, G Danezis 08-2016 USENIX Security

Traffic confirmation attacks despite noise
J Hayes 02-2016 NDSS (workshop track)

Guard sets for onion routing
J Hayes 07-2015 PETS

An introduction to the dynamics of real and complex quadratic polynomials
J Hayes 07-2011

PRESENTATIONS

On visible adversarial perturbations & digital watermarking, CVPR Workshop track

Learning universal adversarial perturbations with generative models, DLS (IEEE S&P Workshop)

Invited talk on Adversarial Machine Learning
IBM, IBM Thomas J. Watson Research Center

Invited talk on Network Traffic Analysis
UK Gov

TASP: Towards Anonymity Sets that Persist
WPES (Workshop of CCS)

k -fingerprinting: a Robust Scalable Website Fingerprinting Technique
USENIX Security

Traffic Confirmation Attacks Despite Noise
Understanding and Enhancing Online Privacy (Workshop of NDSS)

Guard Sets for Onion Routing
PETS

Guard Sets for Onion Routing
University College London Information Security Seminar

Secure Sets in Graphs
Programming, Logic, and Semantics Reading Group, Computer Lab, University of Cambridge

PROGRAM
COMMITTEE
MEMBER

AAAI 2020

CVPR 2020

NeurIPS 2019, 2020

ICLR 2020 Workshop on "Towards Trustworthy ML: Rethinking Security and Privacy for ML"

PrivateNLP'20 (WSDM 2020 Workshop on Privacy and Natural Language Processing)

ICML 2019

PPML19 (Privacy-Preserving Machine Learning - CCS 2019 Workshop)

Privacy Enhancing Technologies Symposium 2018, 2019, 2020, 2021

NeurIPS 2018 Workshop on Security in Machine Learning

(External) CCS, 2017

(External) NDSS, 2016

(External) IEEE Symposium on Security and Privacy, 2016

Transactions on Information Forensics & Security

Transactions on Pattern Analysis and Machine Intelligence

TEACHING

Advisor (along with George Danezis) for Hugo Picq's MSc Thesis, 'Understanding and Scaling a Robust Neural Network', 2019

Advisor for Axel Goetz's MSc Thesis, 'Evaluating the use of Deep Learning for Website Fingerprinting', 2017

SKILLS

Knowledge of:

Python • Go • Shell • \LaTeX • Unix/Linux • Vim • Machine Learning (TensorFlow, PyTorch, JAX, SKLearn)

Exposure to:

C++ • Haskell • JavaScript • C • SQL • HTML • CSS • Git

CODE AND SIDE PROJECTS

Public code available at <https://github.com/jhayes14>, code for private projects available on request.

REFERENCES

Please inform me if references are to be contacted.

George Danezis

Reader in Security and Privacy Engineering (Professor)

Email: g.danezis@ucl.ac.uk

University College London,

Dept. of Computer Science,

Gower Street, London WC1E 6BT, U.K.

Louise Walker

Reader in Mathematics
Email: Louise.Walker[at]manchester.ac.uk
Room 2.243, Alan Turing Building
School of Mathematics,
University of Manchester
Oxford Road, Manchester M13 9PL, UK