

Hal Cooper

<https://ieor.columbia.edu/hal-cooper>

Citizenship: British and Australian (E-3 Visa and STEM-OPT Eligible)

Engaged to a US Citizen. Visa sponsorship not required.

Email: hal.cooper@columbia.edu

Mobile 1: +1-704-449-7832

Mobile 2: +1-646-919-4082

PERSONAL SUMMARY

I'm a doctoral candidate in Operations Research at Columbia, interested in graph computing, machine learning, and interpretable models. I've applied this interest to problems from a variety of fields, including finance, computational neuroscience, sports analytics, and information retrieval. Through close connections with industry (including joint study agreements and internships at Goldman Sachs, IBM, and Graphen, Inc.), I've used state of the art theory to tackle practical problems. I would enjoy employing this same approach to complex problems in an industry role.

EDUCATION

- **Columbia University in the City of New York** New York, NY, USA
MS/MPhil/PhD in Operations Research; GPA: 3.96 *Sep. 2013 – Present*
My advisor in this program is Professor and Department Chair, Garud Iyengar. The program is a joint Masters/Doctoral program, and included an MS in Operations Research, conferred Fall 2014. I was similarly granted an MPhil after my successful dissertation proposal in Fall 2017. My dissertation defense is planned for August 2019.
- **University of Newcastle** Newcastle, NSW, Australia
Bachelor of Engineering in Electrical Engineering; GPA: 7.00 (7-point scale) *Feb. 2008 – Jul. 2012*
I graduated from this program with Honours Class I and the University Medal. As part of this degree, I studied abroad at the University of California, Berkeley, and the University of Virginia.

TECHNICAL

- **Languages:** C++, Python, C, C#, Go, Java, Julia, Scala, SQL, VHDL.
- **Technologies:** Gremlin, L^AT_EX, Linux, Matlab, Neo4j, Spark, Tensorflow, Theano.

DOCTORAL DISSERTATION

- **SmartGraph: An Artificially Intelligent Graph Database**
Graph databases and distributed graph computing systems have traditionally abstracted the design and execution of graph algorithms by encouraging users to take the perspective of static graph objects, like vertices and edges. This dissertation introduces the SmartGraph, a graph database that instead relies upon thinking like a smarter device often found in real-life computer networks, the router. By using a router inspired abstraction to encapsulate vertices and edges, the database facilitates the use of multiple graph representations with local (implementation specific) indexing and highly concurrent execution of analytics. Furthermore, this abstraction allows for artificial intelligence capabilities to be integrated as a centerpiece of the SmartGraph, demonstrating a new generation of graph database technologies.

WORK EXPERIENCE

- **Columbia University** New York, NY, USA
Teaching Assistant and Department Research Assistant *Sep. 2013 - Present*
 - **DRA:** Beginning in Fall 2016, I became a Department Research Assistant. In this position, I had a Joint Study Agreement with IBM from Fall 2016 to Summer 2017. As part of this JSA, I helped IBM submit two DARPA research proposals.

- **TA for Optimization Models and Methods for Financial Engineering:** This course involved teaching methods for modeling problems in financial engineering as optimization problems, as well as the techniques and programming languages (Python, R, Matlab) used to solve such problems. I was the TA for this course twice.
- **TA for Experimental Finance:** This course involved using SQL to investigate real-world financial datasets. I was the TA for this course twice.
- **TA for Advanced Programming for Financial Engineers 1:** This course involved teaching beginner and intermediate Python programming. I was the TA for this course twice.
- **TA for Advanced Programming for Financial Engineers 2:** This course involved teaching advanced programming concepts (in particular, distributed computing and multi-threaded programming in C). I was the TA for this course once.

- **Graphen, Inc.**

New York, NY, USA

- *Summer Intern/Associate*

Summer 2017 (Intern), Summer 2018 (Associate)

- **Software Engineering and Research:** In 2017, members of the SystemG group created the startup Graphen, Inc. Though it is a different company, the work and research conducted for Graphen thus resembles a continuation of the work I began with IBM SystemG.

- **IBM**

Yorktown Heights, NY, USA

- *Summer Intern*

Summer 2015, Summer 2016

- **SystemG:** I worked with IBM SystemG, IBM's research and consulting group for graph computing. As part of this research group, I worked on the aforementioned DARPA research proposals, as well as in-house research problems related to informational retrieval and graph computing. This work was a large influence on the direction of my dissertation.

- **Goldman Sachs**

New York, NY, USA

- *Summer Associate*

Summer 2014

- **Investment Management Division:** I worked in an IMD machine learning research group to solve various machine learning problems in the area of portfolio optimization.

RESEARCH PUBLICATIONS

- **H. Cooper, G. Iyengar, and C. Lin, “Deep Influence Diagrams: An Interpretable and Robust Decision Support System,” in 22nd International Conference on Business Information Systems, 2019. In Press.** Further development of work outlined in “Interpretable Robust Decision Making” below.
- **W. Liu, H. Cooper, M. Oh, P. Chen, S. Yeung, F. Yu, and T. Suzumura, “Learning Graph Topological Features via GAN,” IEEE Access, vol. 7, pp. 21834–21843, 2019.** Further development of work outlined in “Can GAN Learn Topological Features of a Graph?” below. Note the authorship for this paper as listed online is incorrect, and is currently being adjusted to the authorship as listed here.
- **H. Cooper, G. Iyengar, and C. Lin, “Smartgraph: An Artificially Intelligent Graph Database,” in 7th International Conference of Advanced Computer Science & Information Technology, 2019, pp. 63–77.** This paper focuses on using local router artificial intelligence for efficient graph structure selection.
- **H. Cooper, G. Iyengar, and C. Lin, “Personalized Product Recommendation for Interactive Media,” in International Conference on Intelligent Human Systems Integration, 2019.** This work introduces a new method for generating customizable product purchase recommendations from network structured datasets.
- **F. Fagan, M. Haugh, and H. Cooper, “The Advantage of Lefties in One-on-One Sports,” J. Quant. Anal. Sport., 2018.** This article developed new Bayesian models for analyzing latent-variables (such as the left-handed competitive advantage) in datasets exhibiting survival bias.

- **H. Cooper, G. Iyengar, and C. Lin, “Interpretable Robust Decision Making,” in International Conference on Autonomous Agents and Multiagent Systems, 2018.** Decision making algorithms often lack interpretability; they might tell us “what” to do, but not “why” it should be done. This paper develops a decision support system that is highly interpretable and facilitates robust modifications to reflect custom uncertainty and risk aversion profiles.
- **Linderman, G. Mena, H. Cooper, L. Paninski, and J. Cunningham, “Reparameterizing the Birkhoff Polytope for Variational Permutation Inference,” in International Conference on Artificial Intelligence and Statistics, 2018, pp. 1618–1627.** This paper develops a continuous approximation to permutation matrices for use in Bayesian inference.
- **W. Liu, P. Chen, H. Cooper, M. Oh, S. Yeung, and T. Suzumura, “Can GAN Learn Topological Features of a Graph?,” presented in ICML 2017 Workshop on Implicit Models, 2017.** Use of GANs for identifying local and global topological structures in multi-layer graphs.
- **H. J. Cooper, G. C. Goodwin, A. Feuer, and M. G. Cea, “Design of Scenarios for Constrained Stochastic Optimization via Vector Quantization,” in American Control Conference, 2012, pp. 1865–1870.** Method for generating samples for stochastic optimization using vector quantization rather than Monte Carlo approaches. Related to my Honours Thesis as described below.

PATENTS

- **Patent Pending on “System And Method For Providing An Artificially-Intelligent Graph Database.”** I am the First Named Inventor to a provisional application for a patent on the system described in my doctoral dissertation research, filed Oct. 4, 2018 with the United States Patent and Trademark Office.
- **Patent Pending on “System and Method for Providing a Graph Protocol for Forming a Decentralized and Distributed Graph Database.”** I am inventor five to a full patent application developed by Graphen, Inc., filed Dec. 24, 2018 with the United States Patent and Trademark Office.

HONOURS THESIS

- **A Novel Method of Stochastic Control with Application to Finance**

This work involved the development of a new stochastic control methodology that combined model predictive control and dynamic programming through the use of vector quantized scenario design (see Publications) and varied scenario tree branching. The method is demonstrably superior to model predictive control in terms of expected performance, but unlike dynamic programming does not suffer from the curse of dimensionality. This methodology was then applied to the optimization problem of optimal portfolio management through an automated trading system on simulated stocks.

REFERENCES

Available upon request.