

# Chris Olah

---

*"I want to understand things clearly and explain them well."*

## Work Experience

- Oct. 2018 - **OpenAI**, Member of Technical Staff.
  - Led new Clarity team working on neural network interpretability.
- Oct. 2016 - 2018 **Google Brain**, Research Scientist.
- Oct. 2015 - 2016 **Google Brain**, Research Associate.
  - Continued basic research in neural networks.
- May - Oct., 2015 **Google Brain**, Intern.  
Host: Greg Corrado
  - Visualized the 'platonic ideal' of classes according to convolutional neural networks.
  - Developed other novel techniques for visualizing neural networks.
- July - Oct, 2014 **Google Brain**, Intern.  
Host: Jeff Dean
  - Explored the use of interactive media for visualizing neural networks representations.
  - Created the meta-SNE algorithm, which can visualize the space of neural networks.
- July - Sep, 2011 **Xelerance**, Intern.
- July - Nov, 2010 **Environment Canada**, Research Assistant.
- July - Aug, 2009 **University of Toronto**, *Dept. of Forestry*, Research Assistant.

## Honours

- July 2012 **Thiel Fellowship**.  
\$100,000 Fellowship that supports exceptional people under the age of 20 pursue research or start companies.
- July 2010 **AP National Scholar**.  
Graduated high school with six AP (university equivalent) credits.

## Published Security Vulnerabilities

- May 2011 **Unbound DNS Resolver DDOS Vulnerability**.  
CVE-2011-1922 / VU#531342

## Review Service

- 2014 **International Conference on Machine Learning**.
- 2014, 2016 **International Conference on Learning Representations**.
- 2014 **Neural Information Processing Systems Deep Learning Workshop**.

## Writing

- March, 2019 **Activation Atlases**, Shan Carter, Zan Armstrong, Ludwig Schubert, Ian Johnson, & Chris Olah. Distill.
- July, 2018 **Differentiable Image Parameterizations**, Alexander Mordvintsev, Nicola Pezzotti, Ludwig Schubert, & Chris Olah. Distill.
- March, 2018 **The Building Blocks of Interpretability**, Chris Olah, Arvind Satyanarayan, Ian Johnson, Shan Carter, Ludwig Schubert, Katherine Ye, & Alexander Mordvintsev. Distill.  
100+ citations
- Feb 2018 **Is Generator Conditioning Causally Related to GAN Performance?**, Augustus Odena, Jacob Buckman, Catherine Olsson, Tom B Brown, Christopher Olah, Colin Raffel, & Ian Goodfellow. Arxiv preprint.  
3 citations
- Nov, 2017 **Feature Visualization**, Chris Olah, Alexander Mordvintsev, & Ludwig Schubert. Distill.  
100+ citations
- March, 2017 **Research debt**, Christopher Olah & Shan Carter. Distill.  
3 citations
- Feb, 2017 **Changing model behavior at test-time using reinforcement learning**, Augustus Odena, Dieterich Lawson, & Christopher Olah. Arxiv preprint.  
11 citations
- Dec, 2016 **Four Experiments in Handwriting with a Neural Network**, Shan Carter, David Ha, Ian Johnson, & Chris Olah. Distill.  
400,000+ views  
7 citations
- Oct, 2016 **Conditional Image Synthesis With Auxiliary Classifier GANs**, Augustus Odena, Chris Olah, & Jon Shlens. ArXiv preprint.  
500+ citation
- Oct, 2016 **Deconvolution and Checkerboard Artifacts**, Augustus Odena, Vincent Dumoulin, & Chris Olah. Distill.  
100,000+ views  
300+ citation
- Sept 8, 2016 **Attention and Augmented Recurrent Neural Networks**, Chris Olah & Shan Carter. Distill.  
400,000+ views  
36 citations
- 2016 **Concrete Problems in AI Safety**, Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, & Dan Mané.  
300+ citation
- 2015 **TensorFlow: Large-scale machine learning on heterogeneous systems**, Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, Xiaoqiang Zheng. Software available from tensorflow.org.  
10,000+ citations
- Oct. 14, 2015 **Visual Information Theory**, colah.github.io.  
400,000+ views
- Sep. 3, 2015 **Neural Networks, Types, and Functional Programming**, colah.github.io.  
100,000+ views
- Aug. 31, 2015 **Calculus on Computational Graphs: Backpropagation**, colah.github.io.  
400,000+ views  
8 citations
- Aug. 27, 2015 **Understanding LSTM Networks**, colah.github.io.  
3,000,000+ views  
500+ citations

- June 17, 2015 **Inceptionism: Going Deeper into Neural Networks**, Google Research Blog. Alexander Mordvintsev, Christopher Olah, & Mike Tyka.  
1,700,000+ views  
300+ citations
- Jan. 16, 2015 **Visualizing Representations: Deep Learning and Human Beings**, colah.github.io.  
100,000+ views
- Dec. 2014 **Document Embedding with Paragraph Vectors**, NIPS Deep Learning Workshop. Andrew M Dai, Christopher Olah, Quoc V Le, & Greg S Corrado.  
200+ citations
- Dec. 8, 2014 **Groups & Group Convolutions**, colah.github.io.  
40,000+ views
- Oct. 9, 2014 **Visualizing MNIST: An Exploration of Dimensionality Reduction**, colah.github.io.  
400,000+ views
- July 13, 2014 **Understanding Convolutions**, colah.github.io.  
200,000+ views
- July 8, 2014 **Conv Nets: A Modular Perspective**, colah.github.io.  
200,000+ views  
11 citations
- July 7, 2014 **Deep Learning, NLP, and Representations**, colah.github.io.  
200,000+ views
- July 6, 2014 **Fanfiction, Graphs, and PageRank**, colah.github.io.  
50,000+ views
- April 6, 2014 **Neural Networks, Manifolds, and Topology**, colah.github.io.  
500,000+ views  
15 citations
- In Progress **A Weird, Motivated, Intuitive, Introduction to Topology.**,  
github.com/colah/.
- July 16, 2013 **Order Statistics**, colah.ca.  
500+ views
- June 9, 2013 **How My Neural Net Sees Blackboards (Part 2)**, colah.ca.  
500+ views
- May 29, 2013 **I'm Sick and Tired of 3D Printed Guns**, colah.ca.  
250+ views
- May 11, 2013 **How My Neural Net Sees Blackboards**, colah.ca.  
500+ views
- June 17, 2012 **Monads for the Terrified**, colah.ca.
- Feb 10, 2012 **Quantified Hacklab (Part 1)**, colah.ca.
- Nov 6, 2011 **Manipulation of Implicit Functions (With an Eye on CAD)**, colah.ca.  
2,000+ views
- Nov 1, 2011 **Producing Lenses with 3D Printers**, Open Hardware Journal.
- Aug 29, 2011 **Understanding Pascal's Triangle**, colah.ca.  
10,000+ views
- Aug. 11, 2011 **You Already Know Calculus: Differential (One) Forms**, colah.ca.  
250+ views
- Aug. 8, 2011 **The Real 3D Mandelbrot Set**, colah.ca.  
15,000+ views
- July 31, 2011 **You Already Know Calculus: Derivatives**, colah.ca.  
500+ views
- July 16, 2011 **Surface-Oriented CAD, Math, & Telescopes**, colah.ca.  
500+ views
- June 6, 2011 **Alien Mathematics, Numbers, and Polynomial Centric Societies**, colah.ca.  
1,000+ views
- April 18, 2011 **Rethinking Topology (or a Personal Topologodicy)**, colah.ca.  
1,000+ views

- March 28, 2011 **Rethinking Grade School Algebra**, colah.ca.  
2,000+ views
- July 8, 2010 **Towards a Better Notation for Mathematics**, colah.ca.

---

## Talks

- June 23, 2016 **Deep Learning Transparency**, ICML 2016 Workshop on Visualization for Deep Learning. Invited Speaker.
- June 1, 2016 **How Neural Networks Bend Data**, Music, Art, & Machine Intelligence Workshop, Google. Invited Speaker.
- March 18, 2016 **Media and Neural Networks**, Tools for Thought Workshop, Recurse Center. Invited Speaker.
- Feb 26, 2016 **How Neural Networks Bend Data**, DeepDream: The art of neural networks, Gray Area Foundation for the Arts. Invited Speaker.
- Feb 10, 2015 **Neural Networks and the Structure of Data**, Intersections KW.
- Nov 17, 2014 **Why Pattern Recognition is Hard, and Why Deep Neural Networks Help**, Waterloo Computer Science Club.
- Jan. 24, 2014 **Visualizing the Space of Neural Network Hyper-Parameters**, Google.
- Sep 30, 2013 **Smart Kids Are Doing it for Themselves**, Equinox Summit: Learning 2030, Perimeter Institute. Invited Panelist.
- July 12, 2013 **3D Printing For Mathematical Visualization**, Canadian Undergraduate Math Conference.
- Nov. 17, 2012 **Constructive Ways to Build a Better Future**, TEDxYouth@Toronto.
- Oct. 13, 2012 **Multiplicative Calculus For Analyzing Exponential Trends**, Singularity Summit.
- April 18, 2012 **3D Printing & ImplicitCAD**, Noisebridge.
- Nov. 8, 2011 **Open Source 3D Printing: The Printers, Toolchain, & Things**, Greater Toronto Area Linux User Group.
- Oct. 1, 2011 **3D Printing Awesome Things**, SoOnCon.
- Sep. 17, 2011 **Programmatic CAD and its Future**, NYC Maker Faire.

---

## Selected Open Source Participation

- 2018 - **Lucid**, Founder.  
[github.com/tensorflow/lucid](https://github.com/tensorflow/lucid) – *Tools for visualizing neural networks.*
- 2011 - 2013 **ImplicitCAD**, Founder.  
[implicit.herokuapp.com](https://implicit.herokuapp.com) – *A programming language that compiles into 3D objects, written in Haskell*
- Implemented geometry engine (primitives, CSG, etc), interpreter, & GCode generation
- Jan - Feb, 2013 **Printrun**, Contributor.  
[github.com/kliment/Printrun](https://github.com/kliment/Printrun) – *Pure Python 3d printing host software*
- Added safety checks and improved CLI interface

- May - March, 2011 **Printable Vacuum Cleaner**, Author.  
[github.com/colah/Printable-Vacuum-Cleaner](https://github.com/colah/Printable-Vacuum-Cleaner)  
*A Hand Held 3D printable vacuum cleaner!*
- May - Sep., 2011 **surfcad**, Author.  
[github.com/colah/surfcad/](https://github.com/colah/surfcad/) – *Surface-Oriented Programmatic CAD*
- May - Jan., 2011 **ldnsx**, Author.  
[github.com/colah/ldnsx/](https://github.com/colah/ldnsx/) – *A better Python Idns interface*
- March - Sep., 2011 **OpenSCAD**, Contributor.  
[www.openscad.org](http://www.openscad.org) – *Programmatic 3D CAD*
- Implemented syntax highlighting and language extensions
- 2010 - 2011 **Malthus RepRap**, Core Developer.  
[github.com/hacklabto/Hacklab-RepRap](https://github.com/hacklabto/Hacklab-RepRap)  
*An open-source 3D printer striving for self-replication, derived from the Prusa Mendel*
- Redesigned parts to reduce print time and increase ease of assembly
- 2009 - 2010 **Sage**, Contributor.  
[www.sagemath.org](http://www.sagemath.org) – *Open-Source Mathematics Software*
- Added support for exporting 3D visualizations as STLs for 3D printing

---

## Leadership

- 2016 - **Distill Journal**, Editor & Founder.  
*A scientific journal dedicated to clear communication.*
- 2009 - 2014 **hacklab.to**, Member & Director.  
*A hackerspace (community technology space) in Toronto*
- Oversaw management of corporation as a Director (Feb 2012 - Feb 2014).
  - Helped maintain the safety and functionality of the physical space.
- 2012 - 2013 **DIY Bio Toronto**, Co-Organizer.  
*Biology enthusiasts looking to start a biohackerspace*
- Organized several meetups.
  - Started the Molecular Biology of the Cell (Alberts, *et al.*) study group.
- 2012 - 2013 **Toronto Haskell Meetup**, Organizer.  
*Haskell enthusiasts*
- Organized monthly meetups.
- 2011 - 2012 **Toronto 3D Printers**, Organizer.  
*3D Printing Enthusiasts*
- Organized the group while it grew from a handful of people to 40+.

---

## Volunteering

- 2015- **Open Philanthropy Project**, Scientific Adviser.  
*Providing scientific expertise on machine learning & artificial intelligence.*
- 2010 - 2012 **Free Byron**, Court Supporter.  
*Documenting the trial of security researcher Byron Sonne*
- Feb. - Sep., 2010 **Fort York Food Bank**, Volunteer.  
*Distributed food to clients*

---

## Teaching

I've taught five seminar series on neural networks (one online, one at GiveWell, and three at Google). I've also taught many different workshops at hacklab.to on topics ranging from Integral Transforms to  $\LaTeX$ .