



Yarin Gal

Contact Details

Department of Computer Science, University of Oxford, Oxford, UK.

Email, yarin.gal@cs.ox.ac.uk (previously: yg279@cam.ac.uk).

Web, <http://yarin.co/>.

Principal Appointments

2017–present **Associate Professor of Machine Learning**, *University of Oxford Department of Computer Science*, Oxford, UK.

2017–present **Tutorial Fellow**, *Christ Church College*, Oxford, UK.

2017–present **Visiting Researcher**, *Machine Learning Group*, University of Cambridge, UK.

2016–present **Research Fellow**, *Alan Turing Institute for Data Science*, London, UK.

Previous Appointments

2016–2017 **Research Fellow in Computer Science (JRF)**, *St Catharine's College, University of Cambridge*, Cambridge, UK.

Education

2012–2016 **PhD, Information Engineering, Machine Learning Group**, *University of Cambridge*, Cambridge, UK, supervised by Prof Zoubin Ghahramani FRS.
Supported by the Google European Doctoral Fellowship; Qualcomm Innovation Fellowship; Cambridge trusts honorary scholar.

2011–2012 **MSc, Computer Science with focus on Machine Learning**, *University of Oxford*, Oxford, UK, supervised by Prof Phil Blunsom.
Graduated with distinction.

2004–2009 **BSc, double degree in Mathematics and Computer Science**, *The Open University of Israel*, Israel, *Graduated with distinction.*
Ranked in the top 10% of graduates in 2009 (90.8 percentile); Dean's List for the years 2007-2009; Graduated at 18 years old.

Research Interests

Bayesian deep learning • deep learning • approximate Bayesian inference • Gaussian processes • Bayesian modelling • Bayesian non-parametrics • scalable MCMC • generative modelling; Applications including AI safety • ML interpretability • reinforcement learning • active learning • natural language processing • computer vision • medical analysis.

Department of Computer Science, University of Oxford – Oxford, UK

✉ yarin.gal@cs.ox.ac.uk

1/8

Prizes and Scholarships

- 2016–2019 **The Michael and Morven Heller Research Fellowship in Computer Science.**
- 2015 **Alan Turing Institute Travel Award.**
- 2015–2016 **Qualcomm Innovation Fellowship, [press release](#), [media coverage](#).**
- 2015 **BNP 2015 Travel Award.**
- 2014 **Google DeepMind 2014 Travel Award.**
- 2014 **Art of Engineering photo competition.**
- 2014 **NIPS 2014 Travel Award.**
- 2014 **ISBA 2014 Travel Award.**
- 2013 **Tübingen MLSS grant, [press release \(German\)](#).**
- 2012–2015 **Google European Doctoral Fellowship – full PhD scholarship, [press release](#).**
- 2012 **Cambridge overseas trusts – offer of PhD scholarship.**
- 2012 **Oxford-MAN Institute – offer of PhD scholarship.**
- 2012 **Karten Scholarship – an offer of a £2,000 prize.**
- 2004–2009 **Merit based scholarship for high-school students studying for a BSc.**

Conferences and Peer Reviewing Activities

Grant Reviewing for:

- 2017 **Israeli Science Foundation (ISF).**

Workshop Organiser:

- 2016 **First NIPS Workshop on “Bayesian Deep Learning”.**

Program Committee Member:

- 2017 **Workshop on Human Interpretability in Machine Learning (ICML workshop).**
- 2017 **Workshop on Principled Approaches to Deep Learning (ICML workshop).**
- 2016 **The 54th Annual Meeting of the Association for Computational Linguistics (ACL conference, for the area Machine Learning).**
- 2015 **Advances in Approximate Bayesian Inference (NIPS workshop).**
- 2015 **Bayesian Nonparametrics: The Next Generation (NIPS workshop).**
- 2014 **Advances in Variational Inference (NIPS workshop).**

Journal Articles Reviewing for:

- 2017 **American Astronomical Society, The Astrophysical Journal, Letters.**
- 2017 **Nature.**
- 2016 **Journal of the Royal Statistical Society (RSS).**
- 2016 **Journal of Machine Learning Research (JMLR).**
- 2015 **IEEE Transactions on Neural Networks and Learning Systems (IEEE TNNLS).**

Conference Papers Reviewing for:

- 2017 **Neural Information Processing Systems (NIPS).**
- 2017 **International Conference on Machine Learning (ICML).**
- 2017 **International Conference on Learning Representations (ICLR, conference).**
- 2017 **International Conference on Learning Representations (ICLR, workshop).**

- 2016 **Neural Information Processing Systems (NIPS).**
- 2016 **International Conference on Machine Learning (ICML).**
- 2016 **International Conference on Learning Representations (ICLR, conference).**
- 2016 **International Conference on Learning Representations (ICLR, workshop).**
- 2016 **Association for Computational Linguistics (ACL).**
- 2015 **Artificial Intelligence and Statistics (AISTATS).**
- 2015 **Neural Information Processing Systems (NIPS).**
- 2015 **International Conference on Machine Learning (ICML).**

Invited Participant:

- 2017 **UK–Canada AI Research delegation, British Embassy, Montreal.**
- 2017 **UK–Japan AI Research delegation and workshop, British Embassy, Tokyo.**
- 2017 **Google Machine Learning Summit, Zurich.**
- 2016 **GCHQ round-table discussion leader on machine learning, Government Communications Headquarters, UK.**
- 2016 **NASA–SETI machine learning think-tank, NASA's Ames Research Center / SETI Institute.**
- 2016 **NVIDIA GPU machine learning tech summit, Santa Clara.**
- 2015 **First Deep Learning Symposium at NIPS 2015, Montreal.**
- 2015 **Alan Turing Institute Scoping Workshop on Deep Learning, Edinburgh.**
- 2015 **Google NLP PhD Summit, Zurich.**
- 2015 **“Bayesian Nonparametrics in the North” workshop, Ecole Centrale de Lille.**
- 2015 **Google Doctoral Fellowship Forum, Zurich.**
- 2013 **Google Doctoral Summit, Zurich.**

Recent Invited Talks

- 2017 **Preferred Networks, Tokyo, Japan.**
- 2017 **Riken, Tokyo, Japan.**
- 2017 **34th International Conference on Machine Learning, Sydney, Australia.**
- 2017 **Robotics: Science and Systems (RSS) Conference, New Frontiers for Deep Learning in Robotics workshop and panel discussion, Massachusetts, US.**
- 2017 **NASA AMES, Mountain View.**
- 2017 **O’Reilly Artificial Intelligence, New York, US.**
- 2017 **British Embassy, Tokyo, Japan.**
- 2017 **SoftBank, Tokyo, Japan.**
- 2017 **Preferred Networks, Tokyo, Japan.**
- 2017 **Fujitsu, Tokyo, Japan.**
- 2017 **Department of International Trade event, Nagoya, Japan.**
- 2016 **NASA AMES, Mountain View.**
- 2016 **Alan Turing Institute Deep Generative Models Workshop, Royal Society, London.**
- 2016 **Adaptive Brain Lab, Department of Psychology, Cambridge.**

- 2016 **MRC Cognition and Brain Science Unit**, *Cambridge*.
- 2016 **London Machine Learning Meetup**, *London*.
- 2016 **OpenAI**, *San Francisco*.
- 2016 **Google**, *Mountain View*.
- 2016 **33rd International Conference on Machine Learning**, *New York City, NY*.
- 2016 **Natural Language and Information Processing seminar series**, *University of Cambridge, Cambridge*.
- 2015 **Microsoft Research**, *Cambridge*.
- 2015 **Alan Turing Institute Deep Learning Open Workshop**, *Edinburgh University, Edinburgh*.
- 2015 **ATI Scoping Workshop**, *Edinburgh University, Edinburgh*.
- 2015 **ML Seminar series**, *ETH, Zurich*.
- 2015 **Computational Statistics and Machine Learning seminar series**, *UCL, London*.
- 2015 **Bayesian Nonparametrics in the North meeting**, *Ecole Centrale de Lille, Lille*.
- 2015 **Trinity College Mathematical Society**, *University of Cambridge, Cambridge*.
- 2015 **Gonville and Caius College**, *University of Cambridge, Cambridge*.
- 2015 **32nd International Conference on Machine Learning**, *Lille*.
- 2015 **NTT Labs**, *Kyoto, Japan*.
- 2015 **Microsoft Research**, *Cambridge*.
- 2015 **32nd International Conference on Machine Learning**, *Lille*.
- 2015 **10th Conference on Bayesian Nonparametrics**, *Raleigh, NC*.
- 2014 **NTT Labs**, *Kyoto, Japan*.
- 2014 **Workshop on New Learning Models and Frameworks for Big Data, ICML**, *Beijing, China*.
- 2013 **Workshop on Big Learning, NIPS**, *Lake Tahoe*.
- 2013 **Association for Computational Linguistics (NA-AACL)**, *Atlanta*.

Seminar Talks and Reading Groups

- 2016 **“Differentiable Data Structures (and POMDPs)”**, *MLG Seminar*.
- 2014 **“Symbolic Differentiation for Rapid Model Prototyping in Machine Learning and Data Analysis – a Hands-on Tutorial”**, *MLG Seminar*.
- 2014 **“Rapid Prototyping of Probabilistic Models using Stochastic Variational Inference”**, *Short talk*.
- 2014 **“Emergent Communication for Collaborative Reinforcement Learning”**, *MLG Seminar*.
- 2014 **“The Borel–Kolmogorov paradox”**, *Short talk*.
- 2013 **“Bayesian Nonparametrics in Real-World Applications: Statistical Machine Translation and Language Modelling on Big Datasets”**, *MLG Seminar*.

Academic Supervision

- 2016–2017 **Piotr Dabkowski**, *Primary MPhil Thesis Supervision*, Computer Laboratory, University of Cambridge.
- 2016–2017 **Jiří Hron**, *Research Assistant Supervision*, Department of Engineering, University of Cambridge.
- 2015–2016 **Riashat Islam**, *Primary MPhil Thesis Supervision*, Department of Engineering, University of Cambridge.
- 2015–2016 **Jiří Hron**, *Primary MSc Thesis Supervision*, Department of Computer Science, University College London (UCL).
- 2015–2016 **Ambrish Rawat**, *Primary MPhil Thesis Supervision*, Department of Engineering, University of Cambridge.

Select Publications

- [1] Piotr Dabkowski and Yarín Gal. Real Time Image Saliency for Black Box Classifiers. In *arXiv:1705.07857*, May 2017.
- [2] Yarín Gal, Jiri Hron, and Alex Kendall. Concrete Dropout. In *arXiv:1705.07832*, May 2017.
- [3] Alex Kendall, Yarín Gal, and Roberto Cipolla. Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics. In *arXiv:1705.07115*, May 2017.
- [4] Alex Kendall and Yarín Gal. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In *arXiv:1703.04977*, March 2017.
- [5] Yarín Gal, Riashat Islam, and Zoubin Ghahramani. Deep Bayesian Active Learning with Image Data. In *Proceedings of the 34th International Conference on Machine Learning (ICML-17)*, 2017.
- [6] Yingzhen Li and Yarín Gal. Dropout Inference in Bayesian Neural Networks with Alpha-divergences. In *Proceedings of the 34th International Conference on Machine Learning (ICML-17)*, 2017.
- [7] Yarín Gal and Zoubin Ghahramani. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in Neural Information Processing Systems 29 (NIPS)*, 2016.
- [8] Yarín Gal, Rowan McAllister, and Carl E. Rasmussen. Improving PILCO with Bayesian neural network dynamics models. In *Data-Efficient Machine Learning workshop, ICML*, April 2016.
- [9] Yarín Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on Machine Learning (ICML-16)*, 2016.
- [10] Yarín Gal and Zoubin Ghahramani. On modern deep learning and variational inference. In *Advances in Approximate Bayesian Inference workshop, NIPS*, 2015.
- [11] Yarín Gal. Rapid prototyping of probabilistic models: Emerging challenges in variational inference. In *Advances in Approximate Bayesian Inference workshop, NIPS*, 2015.

- [12] Yarin Gal and Zoubin Ghahramani. Bayesian convolutional neural networks with Bernoulli approximate variational inference. In *4th International Conference on Learning Representations (ICLR) workshop track*, 2016.
- [13] Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Insights and applications. In *Deep Learning Workshop, ICML*, 2015.
- [14] Yarin Gal and Richard Turner. Improving the Gaussian process sparse spectrum approximation by representing uncertainty in frequency inputs. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, 2015.
- [15] Yarin Gal, Yutian Chen, and Zoubin Ghahramani. Latent Gaussian processes for distribution estimation of multivariate categorical data. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, 2015.
- [16] Hong Ge, Yarin Gal, and Zoubin Ghahramani. Dirichlet fragmentation processes. In *arXiv:1509.04781*, 2015.
- [17] Yarin Gal, Tomoharu Iwata, and Zoubin Ghahramani. An infinite product of sparse Chinese restaurant processes. In *10th Conference on Bayesian Nonparametrics (BNP)*, 2015.
- [18] Yarin Gal, Yutian Chen, and Zoubin Ghahramani. Latent Gaussian processes for distribution estimation of multivariate categorical data. In *Workshop on Advances in Variational Inference, NIPS*, 2014.
- [19] Yarin Gal, Mark van der Wilk, and Carl Rasmussen. Distributed variational inference in sparse Gaussian process regression and latent variable models. In *Advances in Neural Information Processing Systems 27 (NIPS)*. 2014.
- [20] Yarin Gal, Mark van der Wilk, and Carl E. Rasmussen. Distributed variational inference in sparse Gaussian process regression and latent variable models. In *Workshop on New Learning Models and Frameworks for Big Data, ICML*, 2014.
- [21] Yarin Gal and Zoubin Ghahramani. Feature partitions and multi-view clusterings. International Society for Bayesian Analysis (ISBA), 2014.
- [22] Yarin Gal and Zoubin Ghahramani. Pitfalls in the use of parallel inference for the Dirichlet process. In *Proceedings of the 31th International Conference on Machine Learning (ICML-14)*, 2014.
- [23] Yarin Gal and Zoubin Ghahramani. Pitfalls in the use of parallel inference for the Dirichlet process. In *Workshop on Big Learning, NIPS*, 2013.
- [24] Yarin Gal and Mark van der Wilk. Variational inference in the Gaussian process latent variable model and sparse GP regression – a gentle tutorial. *arXiv:1402.1412*, 2014.
- [25] Yarin Gal. Semantics, modelling, and the problem of representation of meaning – a brief survey of recent literature. Technical report, University of Cambridge, 2013.
- [26] Yarin Gal and Phil Blunsom. A systematic Bayesian treatment of the IBM alignment models. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Atlanta, Georgia, June 2013. Association for Computational Linguistics.

- [27] Yarin Gal. Relaxing HMM alignment model assumptions for machine translation using a Bayesian approach. Master's thesis, University of Oxford, 2012.
- [28] Yarin Gal and Mireille Avigal. Overcoming Alpha-Beta limitations using evolved artificial neural networks. In *Ninth International Conference on Machine Learning and Applications (ICMLA)*. IEEE, 2010.

Open Source Activities

GitHub, <https://github.com/yaringal>.

Contributor:

- 2015–2016 **Keras**, *Deep learning package*, Contributions including: Transposed convolutions, RNN dropout, Deconvolutional VAE example.

Industrial Experience

Steering Committees:

- 2017 **NASA Frontier Development Lab**, *Machine learning incubator for space mission projects*.

Consultancies:

- 2017 **NASA**, *Data Science Coach*, NASA Frontier Development Lab.
- 2017 **blank.ai**, *Machine learning advisory*.
- 2016 **NASA**, *Data Science Coach*, NASA Frontier Development Lab.

Vocational

- 2015 **Google Zurich**, *Research Engineer Intern*, Natural Language Processing group, 08/2015 to 01/2016.
- 2008–2011 **Software Engineer, head of Mobile Platforms development**, *IDesia Biometrics*, Caesarea, Israel.
A software company developing ECG based biometrics and consumer healthcare solutions.
Projects:
 - UX and UI logic design;
 - SDK development and code porting for mobile platforms:
 - MTK (Nucleus based embedded OS), Linux (Android, Maemo), Symbian, Windows mobile;
 - QA tools development, offline simulation tools for the National Physical Laboratory (NPL).
- 2006–2007 **Web development**, *The TAO Centre*.
Commercial ASP web programming.

Volunteer work

- 2014-2016 **Member of Gonville and Caius College MCR committee**, *Computing officer*, Designed and deployed MCR website.
- 2013-2014 **Member of Gonville and Caius College MCR committee**, *Dining officer*, Initiated MCR dinners at Caius.
- 2007 **Teaching computer skills to the elderly**.

Interests

- computing Computational aspects of pretty proofs, reverse engineering (windows platforms).

Department of Computer Science, University of Oxford – Oxford, UK

✉ yarin.gal@cs.ox.ac.uk

7/8

cool projects [extrapolated art](#) – winner of the Art of Engineering photo competition (2nd prize; media coverage: [IFLScience](#), [The Telegraph](#), [Daily Mail](#), [Wolfram blog](#), [Hacker News](#), [Kottke](#), [Boing Boing](#) (by [Cory Doctorow](#)), [Gigazine](#), [Habrahabr](#)).

Origami model design, autonomous game playing (2009), automated arbitrage search (2007), Java based database system (2006), web controlled robotic camera (2004).

Languages

English **Excellent**
Hebrew **Native**
Japanese **Studying**