Yuge Shi

Machine Learning PhD Student,
University of Oxford

w : yugeten.github.io
e : yshi@robots.ox.ac.uk

· Representation learning · Generative models
· Disentanglement · Multimodal learning

Education

DPhil in Engineering Science

University of Oxford, 2018-2022

- **Advised by:**
  - Professor Philip Torr,
  - Professor Siddharth Narayanaswamy.
- **Research area:** Representation learning, self-supervised learning, generative models, robustness, multimodal data.

Bachelor of Engineering (Honours)

Australian National University, 2014-2018

- **Majors:** Electronics and Mechatronics.
- **GPA**: 6.5/7, First Class Honours.

Experience

Research Intern, FAIR New York

Remote, 2020

- **Advised by:** Gabriel Synnaeve
- **Topic:** Domain generalisation
- **Publication:** Gradient Matching for Domain Generalisation (arXiv 2021).

Research intern, ANU

Canberra, Australia, 2020

- **Advised by:** Professor Richard Hartley
- **Topic:** Action recognition
- **Publication:** Action Anticipation with RBF Kernalized Feature Mapping RNN (ECCV, 2018)

Student Scholar, CSIRO

Sydney, Australia, 2020

- **Advised by:** Zhidong Li, Jianlong Zhou
- **Topic:** Interpretable AI, Applied ML

References

References can be provided upon request.

Publications

2021 Yuge Shi, Jeffrey Seely, Philip H.S. Torr, N. Siddharth, Awni Hannun, Nicolas Usunier, Gabriel Synnaeve.
Gradient Matching for Domain Generalisation.
The Tenth International Conference on Learning Representations, under review. (ICLR 2022)

2021 Tom Joy, Yuge Shi, Philip HS Torr, Tom Rainforth, Sebastian M Schmon, N Siddharth.
Learning Multimodal VAEs through Mutual Supervision
The Tenth International Conference on Learning Representations, under review. (ICLR 2022)

2020 Yuge Shi, Brooks Paige, Philip H.S. Torr, N. Siddharth
Relating by Contrasting: A Data-efficient Framework for Multimodal Generative Models
The Ninth International Conference on Learning Representations, 2021. (ICLR 2021)

2019 Yuge Shi*, N. Siddharth*, Brooks Paige, Philip H.S. Torr
Variational mixture-of-experts autoencoders for multimodal deep generative models
Advances in Neural Information Processing Systems, vol. 32, 2019, pp. 1569215703. (NeurIPS 2019)

2018 Yuge Shi ,Basura Fernando ,Richard I. Hartley
Action Anticipation with RBF Kernelized Feature Mapping RNN
Proceedings of the European Conference on Computer Vision, 2018, pp. 305322. (ECCV 2018)

Blog Posts

2021 **An incomplete and slightly outdated literature review on augmentation based self-supervise learning.**
https://yugeten.github.io/posts/2021/12/ssl/

2020 **How I learned to stop worrying and write ELBO (and its gradients) in a billion ways.**
https://yugeten.github.io/posts/2020/06/elbo/

2019 **Gaussian Process, not quite for dummies.** *The Gradient.*
https://yugeten.github.io/posts/2019/09/GP/

Talks

Gradient Matching for Domain Generalisation
Multidisciplinary University Research Initiative, *Dec 2021*

Representation learning from multimodal data
AMLab, University of Amsterdam, *Feb 2021*

How I learned to stop worrying and write ELBO in a billion ways
Bayesian Exploration Lab, University of Oxford, *Nov 2020*

Services & Outreach

Co-founder	**GirlsWhoML** *A volunteer-based initiative that aims at providing free education and mentorship in Machine Learning for people who identify as female/non-binary.*
Committee	**Oxford Women in Computer Science Society** *I work closely with various industry sponsors (Google, Github) of OxWoCS; I also organise the monthly seminar workshops.*
Organiser	**Bad Hypothesis Contest (2021 ICLR social)** *We double down on correlation indicates causation and invite speakers to talk about ludicrous but hilarious "scientific" ideas.*
Reviewer	**NeurIPS, ICLR, ICML, IEEE**