

Yingzhen Li

Department of Computing, Imperial College London
South Kensington Campus, London, SW7 2AZ, UK

Email: [yingzhen.li \[at\] imperial.ac.uk](mailto:yingzhen.li@imperial.ac.uk)
Homepage: <http://www.yingzhenli.net>

PROFESSIONAL EXPERIENCE

Imperial College London

Senior Lecturer (equiv. US associate professor)

Lecturer (equiv. US assistant professor)

Department of Computing

London, UK

Sept 2023 -

Jan 2021 - Aug 2023

- ◇ Research interests: probabilistic machine learning, uncertainty quantification, deep generative models, robust AI/ML, sequence modelling, causal machine learning.

The Alan Turing Institute

Turing Fellow

London, UK

Mar 2024 -

Microsoft Research

Senior Researcher

Machine Intelligence Group, All Data AI Group

Cambridge, UK

Aug 2018 - Dec 2020

- ◇ Responsibilities: develop high-impact research agenda in machine learning; collaborate with researchers and product teams to develop new AI products; supervise post-doc researchers, interns and AI residents; write research code and publish research papers at top-tier conferences.

Disney Research

Research Intern

Mentor: Dr. Stephan Mandt

Pittsburgh, PA, USA

June - Sept 2017

- ◇ Developed a generative system for video/audio synthesis that learned disentangled representation of object/speaker identity and pose/speech content without labels or supervisions.
- ◇ Published an ICML 2018 paper, and filed a patent application (see publications).

EDUCATION

University of Cambridge

Ph.D. in Engineering

Machine Learning Group, supervisor: Prof. Richard E. Turner

Cambridge, UK

July 2018

- ◇ Dissertation: *Approximate Inference: New Visions*
Committee: Prof. Zoubin Ghahramani & Dr. Sebastian Nowozin

Sun Yat-sen University

B.S. in Mathematics

Guangzhou, China

June 2013

- ◇ Graduated with distinction (5%), GPA overall 3.9/4.0, major 4.0/4.0

PUBLICATIONS

Google Scholar page: <https://scholar.google.com/citations?user=gcfs8N8AAAAJ&hl=en>

CONFERENCE PUBLICATIONS (*★ representative papers*)

Hee Suk Yoon, Eunseop Yoon, Joshua Tian Jin Tee, Mark A. Hasegawa-Johnson, **Yingzhen Li** and Chang D. Yoo. C-TPT: Calibrated Test-Time Prompt Tuning for Vision-Language Models via Text Feature Dispersion. *International Conference on Learning Representations (ICLR)*, 2024.

Tobias Schröder, Zijing Ou, Jen Ning Lim, **Yingzhen Li**, Sebastian Vollmer and Andrew Duncan. Energy Discrepancies: A Score-Independent Loss for Energy-Based Models. *Neural Information Processing Systems (NeurIPS)*, 2023.

Harrison Zhu, Carles Balsells Rodas and **Yingzhen Li**. Markovian Gaussian Process Variational Autoencoders. *International Conference on Machine Learning (ICML)*, 2023.

Wenlong Chen and **Yingzhen Li**. Calibrating Transformers via Sparse Gaussian Processes. *International Conference on Learning Representations (ICLR)*, 2023.

Hee Suk Yoon, Joshua Tian Jin Tee, Gwangsu Kim, Eunseop Yoon, Sunjae Yoon, **Yingzhen Li** and Chang D. Yoo. ESD: Expected Squared Difference as a Tuning-Free Trainable Calibration Measure. *International Conference on Learning Representations (ICLR)*, 2023.

Ryutaro Tanno, Melanie F. Pradier, Aditya Nori and **Yingzhen Li**. Fixing Neural Networks by Leaving the Right Past Behind. *Neural Information Processing Systems (NeurIPS)*, 2022.

Zijing Ou, Tingyang Xu, Qinliang Su, **Yingzhen Li**, Peilin Zhao and Yatao Bian. Learning Neural Set Functions Under the Optimal Subset Oracle. *Neural Information Processing Systems (NeurIPS)*, 2022 (long oral, 1%).

Yanzhi Chen, Weihao Sun, **Yingzhen Li** and Adrian Weller. Scalable Infomin Learning. *Neural Information Processing Systems (NeurIPS)*, 2022.

Hippolyt Ritter, Martin Kukla, Cheng Zhang and **Yingzhen Li**. Sparse Uncertainty Representation in Deep Learning with Inducing Weights. *Neural Information Processing Systems (NeurIPS)*, 2021.

Thomas Henn, Yasukazu Sakamoto, Clément Jacquet, Shunsuke Yoshizawa, Masamichi Andou, Stephen Tchen, Ryosuke Saga, Hiroyuki Ishihara, Katsuhiko Shimizu, **Yingzhen Li** and Ryutaro Tanno. A Principled Approach to Failure Analysis and Model Repairment: Demonstration in Medical Imaging. *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2021.

Wenbo Gong, Kaibo Zhang, **Yingzhen Li** and José Miguel Hernández-Lobato. Active Slices for Sliced Stein Discrepancy. *International Conference on Machine Learning (ICML)*, 2021.

Wenbo Gong, **Yingzhen Li** and José Miguel Hernández-Lobato. Sliced Kernelized Stein Discrepancy. *International Conference on Learning Representations (ICLR)*, 2021.

Ruqi Zhang, **Yingzhen Li**, Chris De Sa, Sam Devlin and Cheng Zhang. Meta-Learning for Variational Inference. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.

Yi Zhu, Ehsan Shareghi, **Yingzhen Li**, Roi Reichart and Anna Korhonen. Combining Deep Generative Models and Multi-lingual Pretraining for Semi-supervised Document Classification. *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, 2021.

Andrew Y. K. Foong*, David R. Burt*, **Yingzhen Li** and Richard E. Turner. On the Expressiveness of Approximate Inference in Bayesian Neural Networks. *Neural Information Processing Systems (NeurIPS)*, 2020.

Cheng Zhang, Kun Zhang and **Yingzhen Li**. A Causal View on Robustness of Neural Networks. *Neural Information Processing Systems (NeurIPS)*, 2020.

Maximilian Igl, Kamil Ciosek, **Yingzhen Li**, Sebastian Tschiatschek, Cheng Zhang, Sam Devlin and Katja Hofmann. Generalization in Reinforcement Learning with Selective Noise Injection and Information Bottleneck. *Neural Information Processing Systems (NeurIPS)*, 2019.

- ★ **Yingzhen Li**, John Bradshaw and Yash Sharma. Are Generative Classifiers More Robust to Adversarial Attacks? *International Conference on Machine Learning (ICML)*, 2019.

Chao Ma, **Yingzhen Li** and José Miguel Hernández-Lobato. Variational Implicit Processes. *International Conference on Machine Learning (ICML)*, 2019. (long oral, 4.5%)

Ehsan Shareghi, **Yingzhen Li**, Yi Zhu, Roi Reichart and Anna Korhonen. Bayesian Learning for Neural Dependency Parsing. *Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, 2019.

Wenbo Gong*, **Yingzhen Li*** and José Miguel Hernández-Lobato. Meta-learning for Stochastic Gradient MCMC. *International Conference on Learning Representations (ICLR)*, 2019. (joint first author)

- ★ **Yingzhen Li** and Stephan Mandt. Disentangled Sequential Autoencoder. *International Conference on Machine Learning (ICML)*, 2018.

- ★ **Yingzhen Li** and Richard E. Turner. Gradient Estimators for Implicit Models. *International Conference on Learning Representations (ICLR)*, 2018.

Cuong V. Nguyen, **Yingzhen Li**, Thang D. Bui and Richard E. Turner. Variational Continual Learning. *International Conference on Learning Representations (ICLR)*, 2018.

Yingzhen Li and Yarin Gal. Dropout Inference in Bayesian Neural Networks with Alpha-divergences. *International Conference on Machine Learning (ICML)*, 2017.

- ★ **Yingzhen Li** and Richard E. Turner. Rényi Divergence Variational Inference. *Neural Information Processing Systems (NIPS)*, 2016.

José Miguel Hernández-Lobato*, **Yingzhen Li***, Mark Rowland, Daniel Hernández-Lobato, Thang D. Bui and Richard E. Turner. Black-box Alpha-divergence Minimization. *International Conference on Machine Learning (ICML)*, 2016. (joint first author)

Thang D. Bui, Daneil Hernández-Lobato, José Miguel Hernández-Lobato, **Yingzhen Li** and Richard E. Turner. Deep Gaussian Processes for Regression using Approximate Expectation Propagation. *International Conference on Machine Learning (ICML)*, 2016.

- ★ **Yingzhen Li**, José Miguel Hernández-Lobato and Richard E. Turner. Stochastic Expectation Propagation. *Neural Information Processing Systems (NIPS)*, 2015. (spotlight, 4.5%)

PATENTS

Stephan Mandt and **Yingzhen Li**. Efficient Encoding and Decoding Sequences using Variational Autoencoders. *United States Patent Application 20190392302 A1*. Assignee: Disney Enterprises Inc.

INVITED TALKS

Towards causal deep generative models for sequential data Feb 2024
Invited talk, School of Informatics, University of Edinburgh, Edinburgh, UK

Calibrating Transformers via Sparse Gaussian Processes Feb 2024
ELLIS Manchester seminar (with Wenlong Chen), Manchester, UK

Towards causal deep generative models for sequential data Oct 2023
CHAI Structured learning workshop, Gothenburg, Sweden

Towards causal deep generative models for sequential data Aug 2023
Colloquiums at Tsinghua, Peking, Zhejiang and Fudan Universities, Shanghai AI Lab, Zhejiang Lab

Bayesian deep learning via function-space posterior inference Invited session talk, European meeting of statisticians (EMS 2023), Warsaw, Poland	July 2023
Making Stein’s method great Again for generative modelling? Invited talk, Distance-based methods for machine learning workshop, University College London	June 2023
Sequential generative models Invited lecture, Generative modeling summer school, Copenhagen, Denmark	June 2023
Introduction to Bayesian neural networks Invited lecture, Nordic Probabilistic AI School, 2023	June 2023
Towards causal deep generative models for sequential data Colloquium on “When Causal Infrence Meets Statistical Analysis”, Paris, France	Apr 2023
Variational Bayes for deep learning: Towards function-space UQ Bayes Comp 2023 conference, Levi, Finland	Mar 2023
Robust and adaptive deep learning via Bayesian principles New Faculty Highlights, AAI 2023	Feb 2023
Towards better models beyond Just good accuracies Invited talk, Microsoft Health Futures, Microsoft Research, Cambridge, UK	Dec 2022
Understanding masked Pre-training: fundamentally different? Invited talk, Generative Models and Uncertainty Quantification workshop, Copenhagen, Denmark	Sept 2022
A function-space view tour for Bayesian neural networks Invited lecture, Gaussian Process Summer School 2022, U Sheffield	Sept 2022
Introduction to Bayesian neural networks Invited lecture, Nordic Probabilistic AI School, 2022	June 2022
Evaluating approximate inference for BNNs Invited talk, University of Cape Town, South Africa	Apr 2022
Evaluating approximate inference for BNNs Invited talks at Facebook, ByteDance, Huawei	Sept 2021
Inference with scores: slices, diffusions and flows Invited talk, INNF+ workshop, ICML 2021	July 2021
Evaluating approximate inference for BNNs Invited talk, Women in ML Unworkshop, ICML 2021	July 2021
EBM inference & learning: a personal story Invited talk, Energy-based model workshop, ICLR 2021	May 2021
From parametric models to Gaussian processes UnderstandingDL talk series, Data @ University de Sao Paolo	Apr 2021
Advances in approximate inference Invited tutorial, NeurIPS 2020 (with Cheng Zhang)	Dec 2020
On estimating epistemic uncertainty in deep learning Invited talk, INFORMS annual meeting 2020 (session: uncertainty representations in RL)	Nov 2020
Deep probabilistic modelling for reliable ML systems Colloquiums at UCL, Imperial, U Edinburgh	Mar 2020
On estimating epistemic uncertainty Invited talk, Bayesian deep learning workshop, NeurIPS 2019	Dec 2019

On the uncertainty estimation of Bayesian neural networks Invited talk, UNSURE workshop, MICCAI 2019	Oct 2019
Bayesian neural networks: a function space view tour Invited lecture, Gaussian Process Summer School 2019, U Sheffield	Sept 2019
Gradient estimators of implicit models using Stein’s method Invited talk, Stein’s method workshop, ICML 2019	June 2019
On KL divergence and beyond Invited talk, Language technology seminar, U Cambridge	Nov 2018
Meta-learning for SG-MCMC Invited talk, Uncertainty in deep learning workshop, UAI 2018	Aug 2018
Efficient computation for Bayesian deep learning Colloquiums at MSR Cambridge, U Oxford, OpenAI	Mar 2018
Gradient estimators for implicit models Invited talk, Advances in approximate Bayesian inference workshop, NIPS 2017	Dec 2017
Wild approximate inference: why and how Invited talk, CSML seminar, UCL	Dec 2017
Adversarial attacks and defences Seminar, AI safety reading group, U Cambridge	Nov 2017
Approximate inference with Amortised MCMC Invited talk, CamAIML 2017, MSR Cambridge	Mar 2017
Objective functions for variational auto-encoders Invited talk, Twitter Cortex London	May 2016
Variational inference with Rényi divergence Invited talk, MSR Cambridge	Mar 2016

ACADEMIC SERVICES

JOURNAL REVIEWING

Transactions of Machine Learning Research (TMLR, Action Editor)
 Journal of Machine Learning Research (JMLR)
 IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)
 Artificial Intelligence (AIJ)
 Neural Computing
 Bayesian Analysis
 The Annals of Statistics (AOS)

CONFERENCE REVIEWING

Area chair/Senior PC (equiv., journal action editor):
 AAAI Conference on Artificial Intelligence (AAAI 2020)
 Neural Information Processing Systems (NeurIPS 2019, 2021 - 2024)
 NeurIPS 2022 competition track
 International Conference on Machine Learning (ICML 2019, 2020, 2023, 2024)
 International Conference on Artificial Intelligence and Statistics (AISTATS 2019, 2022, 2023)
 International Conference on Learning Representations (ICLR 2021, 2023, 2024)

PC member (equiv., journal reviewer):

Neural Information Processing Systems (NIPS 2016, 2017, NeurIPS 2018, 2020)
 International Conference on Machine Learning (ICML 2017, 2018)
 International Conference on Learning Representations (ICLR 2017, 2018, 2019, 2020)
 International Joint Conference on Artificial Intelligence (IJCAI 2017, 2019)
 AAAI Conference on Artificial Intelligence (AAAI 2018)
 International Conference on Artificial Intelligence and Statistics (AISTATS 2018, 2020)
 Uncertainty in Artificial Intelligence (UAI 2020)

ORGANISING

Approximate inference in Bayesian deep learning, NeurIPS 2021 competition
 Neural compression workshop, ICLR 2021
 Bayesian deep learning NeurIPS workshop/meetup, NeurIPS 2020/2021
 Deep generative models and downstream applications workshop, NeurIPS 2021
 Symposium on advances in approximate Bayesian inference (AABI), 2020/2021/2022/2023
 Uncertainty quantification for computer vision workshop, ECCV 2022
 Score-based methods workshop, NeurIPS 2022
 Challenges and perspectives in deep generative modelling workshop, Dagstuhl seminar, 2023
 International Conference on Artificial Intelligence and Statistics (AISTATS 2024), Program Chair

MENTORSHIP

IMPERIAL COLLEGE LONDON

PhD students:

On-going: Jacob Yoke Hong Si (Jan 2024 -), Naoki Kiyohara (Computing, Oct 2023 -, with Edward Johns), Isuru Shavindra Jayasekera (Math, Oct 2023 -), Lapo Rastrelli (Computing, Oct 2023 -), Francisco Sumba Toral (Computing, Feb 2023 -), Zijing Ou (Computing, Oct 2022 -), Carles Balsells Rodas (Computing, Oct 2021 -), Wenlong Chen (Computing, Oct 2021 -)

Finished: Harrison Bohua Zhu (Math, Oct 2021 - Sept 2023, with Seth Flaxman)

MRes research students:

Xiaoyao Qiu (Oct 2023 -), Ze Zhang (Oct 2022 - Sept 2023), Pengyang Xie (Oct 2022 - Sept 2023), Jiaming Zhang (Nov 2022 - Sept 2023), Lapo Rastrelli (Oct 2021 - Sept 2022)

MSc/Mphil/MEng project students:

Andreas Makris (Computing, 2024), Moritz Hauschulz (Computing, 2024), Alberto Almagro Sanchez (Computing, 2024), Joseph Aghoghovbia (Computing, 2024), Frederick Sligo-Young (Computing, 2023), Jonathan Hau (Computing, 2023), Jason Lee (Computing, 2023), Alexander Pondaven (EEE, 2023), Yasir Abdi (Earth Science, 2022), Wuhao Chen (Computing, 2022), Yiran Wu (Computing, 2022), Mun Fai Chan (Computing, 2022), Xuesi Yang (Computing, 2022), Qing Pan (Computing, 2022), Gengjian Hu (Computing, 2022), Sylvie Shi (Computing, 2021), Yehui Tang (Computing, 2021), Yiming Zhang (Computing, 2021), Stefan Arisatya Tionanda (Computing, 2021), Yuhao Tang (Computing, 2021), Roxane Fischer (Computing, 2021)

Undergraduate UROP:

Rutwij Patel (Computing, 2023), Zhuoyue Huang (Math, 2022), Fengzhe Zhang (Math, 2022)

MSR CAMBRIDGE INTERN/AI RESIDENT/VISITOR (CO-)SUPERVISION

Research interns:

Maximillian Igl (University of Oxford, Feb - May 2019) Chaochao Lu (University of Cambridge, Feb - May 2019) Ruqi Zhang (Cornell University, June - Aug 2019) Sebastian Lunz (University of Cambridge, June - Aug 2019) Xun Zheng (Carnegie Mellon University, June - Aug 2019) Jooyeon Kim (KAIST, June

- Aug 2019) Haiyan Yin (Nanyang Technological University, Sept - Nov 2019) Marco Federici (University of Amsterdam, Oct 2019 - Feb 2020) Hippolyt Ritter (University College London, June - Aug 2020) Philip Ball (University of Oxford, June - Aug 2020) James Jordon (University of Oxford, July - Sept 2020)

AI residents:

Hiske Overweg (Oct 2018 - Mar 2019) Anna-Lena Popkes (Oct 2018 - Mar 2019) Angus Lamb (Oct 2019 - Mar 2020) Evgeny Saveliev (Oct 2019 - Mar 2020) Sarah Lewis (Oct 2020 - Dec 2020) Tatiana Matejovicova (Oct 2020 - Dec 2020)

Visitors:

Rika Antonova (KTH, Nov 2018 - Feb 2019)

UNIVERSITY OF CAMBRIDGE STUDENT (CO-)SUPERVISION AND COLLABORATION

PhD students:

Chao Ma (Engineering, 2018 - 2019), Wenbo Gong (Engineering, 2018 - 2021), Andrew F. K. Foong (Engineering, 2019 - 2020), David R. Burt (Engineering, 2019 - 2020), Yi Zhu (Linguistics, 2019 - 2021), Victor Prokhorov (Linguistics, 2019 - 2021), Yanzhi Chen (Engineering, 2022 -)

Mphil/MEng project students:

Vera Gangeskar Johne (Engineering, 2016), Alexander Yakub (Engineering, 2016), Wenbo Gong (Engineering, 2017), Andy Renqiao Zhang (Engineering, 2017), Siddharth Swaroop (Engineering, 2017)

Undergraduate project students:

Artem Vasenin (Computer Science, 2017)

PHD COMMITTEE

Steindór Saemundsson (Imperial College London, 2021), Arsenii Ashukha (Higher School of Economics, Russia, 2022), Xiaoyu Bie (INRIA Grenoble Rhone-Alpes, 2023), Shikun Liu (Imperial College London, 2024)

TEACHING

Deep Learning, Imperial College London	Jan - Mar 2024
MRes AI & ML, degree program director	Sept 2023 -
Deep Learning, Imperial College London	Jan - Mar 2023
Mathematics for Machine Learning, Imperial College London	Oct - Dec 2022
Deep Learning, Imperial College London	Jan - Mar 2022
Mathematics for Machine Learning, Imperial College London	Oct - Dec 2021
Deep Learning, Imperial College London	Jan - Mar 2021
Microsoft Research AI residency program (lectures and project mentoring), MSR Cambridge	2018 - 2020
Part IIA 3F7 Information Theory and Coding, U Cambridge	Oct - Dec 2016
Mphil in machine learning (modules/lab sessions), U Cambridge	Oct 2015 - Apr 2016
Part IIA 3F6 software engineering, U Cambridge	Jan - Mar 2015
Part IB C++ computing, U Cambridge	Jan - Mar 2014

AWARDS

New Faculty Highlights, AAAI	2023
ICLR 2021 outstanding area chair, ICLR	2021
ICML travel award (\$1,800), IMLS	2017
Neural Information Processing Systems travel award (\$800), Google	2015
FFTF fellowship award (\$50,000/year), Schlumberger Foundation	2014 - 2018
Best Bachelor's thesis award (3%), Sun Yat-sen University	2013
Student scholarship (20%), 21st Machine Learning Summer School, Kyoto University	2012
Excellent student scholarship (5%), Sun Yat-sen University	2010 - 2013

OUTREACH

Demonstration: "Prompt Engineering for image generation". Imperial Late "AI theme" (Mar 2024).

Talk: "Drawing on a bigger canvas". Imperial Computing WiC showcase day (Feb 2022), Imperial Computing taster courses (June 2021/2022).

Talk: "Generative machine learning: beyond TikTok filters". The 63th London International Youth Science Forum (Aug 2022).

MISCELLANEOUS

SKILLS

Software: python (tensorflow/pytorch), matlab, C/C++.

Languages: native/fluent in Chinese Mandarin/Cantonese and English.

EXTRACURRICULAR

Crew member (women team 2), Darwin College Boat Club, Cambridge, UK	2014 - 2015
Co-organiser/participant, 1st/3rd Hackathon for undergraduates, Guangzhou, China	Aug 2011/2012
University Taekwondo team member, Sun Yat-sen University, Guangzhou, China	2010 - 2012
Student journalist & editor, news center, Sun Yat-sen University, Guangzhou, China	2009 - 2011

NAME IN CHINESE CHARACTERS

Surname: 李 (Li)

Given name: 映真 (Yingzhen)

CV last updated: Mar 2024