

Chin-Wei Huang

Machine-learning researcher working across generative modeling, structured architectures, and AI for Science

[Email](#)[Website](#)[Google Scholar](#)[GitHub](#)[X](#)

RESEARCH PROFILE

Machine-learning researcher with expertise in probabilistic generative modeling, structured model architectures, variational inference, and scalable training systems. My current work focuses on physically grounded machine learning for atomistic modeling and electronic structure. A recurring theme in my research is building useful mathematical or physical structure into models, then exploiting that structure for tractable, scalable computation.

PROFESSIONAL EXPERIENCE

Microsoft Research AI for Science — Senior Researcher

Amsterdam · March 2022–present

- Lead machine-learning research on Skala, broadening its scientific scope and capabilities through learned-functional architecture, scalable training, self-consistent finetuning, and model evaluation.
- Created the architecture underlying Skala and developed its theoretical connection to physically meaningful constraints and atomic cluster or multipole structure.
- Built and improved distributed data, training, evaluation, and experiment-tracking systems that enabled substantially larger models and datasets.
- Analyze model limitations and use them to shape data generation, architecture, optimization, and project milestones.
- Previously led diffusion-based reaction-generation research, supervised work on transition-state generation and rate estimation, and co-developed diffusion-model infrastructure reused by MatterGen.

Google — Student Researcher

Montréal · September–December 2019

- Host: Tyler Zhu. Collaborators: Dilip Krishnan and Laurent Dinh.
- Worked on generative models.

Element AI — Research Intern

Montréal · November 2017–May 2019

- Host: Alexandre Lacoste.
- Worked part-time alongside the Ph.D. on expressive and scalable probabilistic models and variational inference, contributing to work including neural autoregressive flows, annealed variational objectives, hierarchical importance-weighted autoencoders, probability distillation, and Kronecker-flow methods.

Mila, Université de Montréal — Research Intern

Montréal · February–August 2017

- Advisors: Laurent Charlin and Aaron Courville.
- Developed Bayesian Hypernetworks and a preliminary workshop version of Neural Autoregressive Flows.
- Explored multi-source latent-representation models for scene and medical-image modeling.

● **Aengin Technology – Data Engineer and Product Development Support**

Remote / Taipei · August–November 2016

● **Institute for Information Industry (IDEAS) – Data Engineer**

Taipei · April–July 2016

- Built internal tools and applications for data analysis.

● **Polytechnique Montréal, Machine Design Lab – Research Intern**

Montréal · January–May 2015

- Advisor: Sofiane Achiche.
- Applied machine-learning and feature-selection methods to paper-pulp quality prediction.

● **Dow Chemical Company – Environment, Health and Safety Intern**

Taoyuan · July–August 2014

● **National Taiwan University, Polymer Nano-Materials Lab – Research Assistant**

Taipei · August 2013–February 2014

- Advisor: Chi-An Dai.
- Assisted research on synthesizing materials for liquid armor.

EDUCATION

● **Ph.D. in Computer Science – Université de Montréal**

2018–2024

- Advisor: Aaron Courville.
- Thesis: *Probability flows in deep learning*.

● **M.S. in Computer Science – Mila, Université de Montréal**

2017–2018

- Fast-tracked into the Ph.D. program.
- Advisor: Aaron Courville. Co-advisor: Laurent Charlin.

● **Ph.D. studies in Mechanical Engineering – Polytechnique Montréal**

2016–2017 · Left the program to pursue computer science and machine learning

- Advisor: Sofiane Achiche. Co-advisor: Maxime Raison.

● **B.S. in Chemical Engineering – National Taiwan University**

2011–2015

PUBLICATIONS AND RESEARCH OUTPUTS

An asterisk denotes shared first authorship.

2026

- 01 **JOURNAL ARTICLE**. Sebastian Ehlert, Jan Hermann, Thijs Vogels, Victor Garcia Satorras, Stephanie Lanius, Marwin Segler, Klaas J. H. Giesbertz, Derk P. Kooi, Kenji Takeda, **Chin-Wei Huang**, Giulia Luise, Rianne van den Berg, Paola Gori-Giorgi, and Amir Karton. "Accurate Chemistry Collection: Coupled cluster atomization energies for broad chemical space." *Scientific Data* 13, 951 (2026). DOI · arXiv

2025

- 01 **PREPRINT**. Giulia Luise*, **Chin-Wei Huang***, Thijs Vogels*, Derk P. Kooi*, Sebastian Ehlert*, Stephanie Lanius, Klaas J. H. Giesbertz, Amir Karton, Deniz Gunceler, Stefano Battaglia, Gregor N. C. Simm, P. Bernát Szabó, Megan Stanley, Wessel P. Bruinsma, Lin Huang, Xinran Wei, José Garrido Torres, Abylay Katbashev, Rodrigo Chavez Zavaleta, Bálint Máté, Sékou-Oumar Kaba, Roberto Sordillo, Yingrong Chen, David B. Williams-Young, Christopher M. Bishop, Jan Hermann, Rianne van den Berg, and Paola Gori-Giorgi. "Accurate and scalable exchange-correlation with deep learning." arXiv:2506.14665 (2025; revised 2026). arXiv · Code

- 02 **JOURNAL ARTICLE.** Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Sasha Shysheya, Jonathan Crabbé, Lixin Sun, Jake Smith, Bichlien Nguyen, Hannes Schulz, Sarah Lewis, **Chin-Wei Huang**, Ziheng Lu, Yichi Zhou, Han Yang, Hongxia Hao, Jielan Li, Ryota Tomioka, and Tian Xie. "A generative model for inorganic materials design." *Nature* 639, 624–632 (2025). [DOI](#) · [arXiv](#) · [Code](#)

2023

- 01 **JOURNAL ARTICLE.** Marloes Arts, Victor Garcia Satorras, **Chin-Wei Huang**, Daniel Zügner, Marco Federici, Cecilia Clementi, Frank Noé, Robert Pinsler, and Rianne van den Berg. "Two for One: Diffusion Models and Force Fields for Coarse-Grained Molecular Dynamics." *Journal of Chemical Theory and Computation* 19(18), 6151–6159 (2023). [DOI](#) · [arXiv](#)

- 02 **THESIS.** **Chin-Wei Huang**. *Probability flows in deep learning*. Doctoral thesis, Université de Montréal (published 2023; degree conferred 2024). [DOI](#)

2022

- 01 **CONFERENCE PAPER.** **Chin-Wei Huang**^{*}, Milad Aghajohari^{*}, Avishek Joey Bose, Prakash Panangaden, and Aaron Courville. "Riemannian Diffusion Models." *Advances in Neural Information Processing Systems* (2022). [arXiv](#)

- 02 **CONFERENCE PAPER.** Shawn Tan, **Chin-Wei Huang**, Alessandro Sordoni, and Aaron Courville. "Learning to Dequantise with Truncated Flows." *International Conference on Learning Representations* (2022).

2021

- 01 **CONFERENCE PAPER.** **Chin-Wei Huang**, Jae Hyun Lim, and Aaron Courville. "A Variational Perspective on Diffusion-Based Generative Models and Score Matching." *Advances in Neural Information Processing Systems* (2021). Spotlight. [arXiv](#)

- 02 **JOURNAL ARTICLE.** Joseph Paul Cohen, Tianshi Cao, Joseph D. Viviano, **Chin-Wei Huang**, Michael Fralick, Marzyeh Ghassemi, Muhammad Mamdani, Russell Greiner, and Yoshua Bengio. "Problems in the deployment of machine-learned models in health care." *CMAJ* 193(35), E1391–E1394 (2021). Invited article. [DOI](#)

- 03 **CONFERENCE PAPER.** **Chin-Wei Huang**, Ricky T. Q. Chen, Christos Tsirigotis, and Aaron Courville. "Convex Potential Flows: Universal Probability Distributions with Optimal Transport and Convex Optimization." *International Conference on Learning Representations* (2021). [arXiv](#)

2020

- 01 **PREPRINT.** Brady Neal, **Chin-Wei Huang**, and Sunand Raghupathi. "RealCause: Realistic Causal Inference Benchmarking." arXiv:2011.15007 (2020).

- 02 **PREPRINT.** Tianshi Cao, **Chin-Wei Huang**, David Yu-Tung Hui, and Joseph Paul Cohen. "A benchmark of medical out of distribution detection." arXiv:2007.04250 (2020).

- 03 **SYMPOSIUM CONTRIBUTION.** Jae Hyun Lim^{*}, **Chin-Wei Huang**^{*}, Aaron Courville, and Christopher Pal. "Bijective-Contrastive Estimation." *Third Symposium on Advances in Approximate Bayesian Inference* (2020). Contributed talk. [OpenReview](#)

- 04 **CONFERENCE PAPER.** Jae Hyun Lim, Aaron Courville, Christopher Pal, and **Chin-Wei Huang**. "AR-DAE: Towards Unbiased Neural Entropy Gradient Estimation." *International Conference on Machine Learning* (2020).

- 05 **WORKSHOP PAPER.** **Chin-Wei Huang**, Laurent Dinh, and Aaron Courville. "Solving ODE with Universal Flows: Approximation Theory for Flow-Based Models." *ICLR Workshop on Integration of Deep Neural Models and Differential Equations* (2020). Contributed talk.

- 06 **PREPRINT.** **Chin-Wei Huang**, Laurent Dinh, and Aaron Courville. "Augmented Normalizing Flows: Bridging the Gap between Generative Flows and Latent Variable Models." arXiv:2002.07101 (2020).

- 07 **CONFERENCE PAPER.** **Chin-Wei Huang**, Ahmed Touati, Pascal Vincent, Gintare Karolina Dziugaite, Alexandre Lacoste, and Aaron Courville. "Stochastic Neural Network with Kronecker Flow." *International Conference on Artificial Intelligence and Statistics* (2020).

2019

- 01 **WORKSHOP PAPER.** **Chin-Wei Huang**, Ahmed Touati, Pascal Vincent, Gintare Karolina Dziugaite, Alexandre Lacoste, and Aaron Courville. "PAC Bayes Bound Minimization via Kronecker Normalizing Flows." *ICML Workshop on Understanding and Improving Generalization in Deep Learning* (2019).
- 02 **CONFERENCE PAPER.** Fan-Yun Sun, Meng Qu, Jordan Hoffmann, **Chin-Wei Huang**, and Jian Tang. "vGraph: A Generative Model for Joint Community Detection and Node Representation Learning." *Advances in Neural Information Processing Systems* (2019).
- 03 **PREPRINT.** Shawn Tan, Yikang Shen, **Chin-Wei Huang**, and Aaron Courville. "Investigating Biases in Textual Entailment Datasets." arXiv:1906.09635 (2019).
- 04 **CONFERENCE PAPER.** **Chin-Wei Huang**^{*}, Faruk Ahmed^{*}, Kundan Kumar, Alexandre Lacoste, and Aaron Courville. "Probability Distillation: A Caveat and Alternatives." *Conference on Uncertainty in Artificial Intelligence* (2019).
- 05 **PREPRINT.** **Chin-Wei Huang** and Aaron Courville. "Note on the Bias and Variance of Variational Inference." arXiv:1906.03708 (2019).
- 06 **CONFERENCE PAPER.** **Chin-Wei Huang**, Kris Sankaran, Eeshan Dhekane, Alexandre Lacoste, and Aaron Courville. "Hierarchical Importance Weighted Autoencoders." *International Conference on Machine Learning* (2019).
arXiv

2018

- 01 **CONFERENCE PAPER.** **Chin-Wei Huang**, Shawn Tan, Alexandre Lacoste, and Aaron Courville. "Improving Explorability in Variational Inference with Annealed Variational Objectives." *Advances in Neural Information Processing Systems* (2018). arXiv
- 02 **CONFERENCE PAPER.** **Chin-Wei Huang**^{*}, David Krueger^{*}, Alexandre Lacoste, and Aaron Courville. "Neural Autoregressive Flows." *International Conference on Machine Learning* (2018). Oral. [Proceedings](#) · [arXiv](#)
- 03 **PREPRINT.** Yikang Shen^{*}, Shawn Tan^{*}, **Chin-Wei Huang**^{*}, and Aaron Courville. "Generating Contradictory, Neutral, and Entailing Sentences." arXiv:1803.02710 (2018).
- 04 **CONFERENCE PAPER.** Yikang Shen, Zhouhan Lin, **Chin-Wei Huang**, and Aaron Courville. "Neural Language Modeling by Jointly Learning Syntax and Lexicon." *International Conference on Learning Representations* (2018).

2017

- 01 **PREPRINT.** **Chin-Wei Huang**, Luc Baron, Marek Balazinski, and Sofiane Achiche. "Comprehensive model optimization in pulp quality prediction: a machine learning approach." *PeerJ Preprints* (2017). [DOI](#)
- 02 **WORKSHOP PAPER.** David Krueger^{*}, **Chin-Wei Huang**^{*}, Riashat Islam, Ryan Turner, Alexandre Lacoste, and Aaron Courville. "Bayesian Hypernetworks." *NeurIPS Workshop on Bayesian Deep Learning* (2017). [arXiv](#)
- 03 **WORKSHOP PAPER.** **Chin-Wei Huang**^{*}, David Krueger^{*}, and Aaron Courville. "Facilitating Multimodality in Normalizing Flows." *NeurIPS Workshop on Bayesian Deep Learning* (2017).
- 04 **WORKSHOP PAPER.** **Chin-Wei Huang** and Aaron Courville. "Sequentialized Sampling Importance Resampling and Scalable IWAE." *NeurIPS Workshop on Bayesian Deep Learning* (2017).

05 **WORKSHOP PAPER.** **Chin-Wei Huang**, Ahmed Touati, Laurent Dinh, Michal Drozdal, Mohammad Havaei, Laurent Charlin, and Aaron Courville. "Learnable Explicit Density for Continuous Latent Space and Variational Inference." *ICML Workshop on Principled Approaches to Deep Learning* (2017).

06 **SYMPOSIUM CONTRIBUTION.** **Chin-Wei Huang** and Nan Rosemary Ke. "Deconstructive Defense Against Adversarial Attacks." *Montréal AI Symposium* (2017).

07 **SYMPOSIUM CONTRIBUTION.** Michal Drozdal, Mohammad Havaei, **Chin-Wei Huang**, Laurent Charlin, Nicolas Chapados, and Aaron Courville. "Data Imputation with Latent Variable Models." *Montréal AI Symposium* (2017).

PATENT

- **Chin-Wei Huang**, inventor. "Denoising diffusion model for coarse-grained molecular dynamics." U.S. Patent Application 18/314,630; publication US20240266004A1. Filed May 9, 2023; published August 8, 2024. [Google Patents](#)

RESEARCH SOFTWARE

sdeflow-light: minimalist implementation of score-based diffusion models. **NAF**: experiments for Neural Autoregressive Flows.

Skala: public inference implementation for the Skala learned density functional.

MatterGen: public implementation of MatterGen; contribution included shared diffusion-model infrastructure and technical advising.

INVITED TALKS AND LECTURES

2025

- "Skala: Accurate and scalable exchange-correlation with deep learning." Molecular Machine Learning Conference (MoML @ MIT), Cambridge, October 22.
- "Accurate and scalable density functional with deep learning." Microsoft Research New England Generative Modeling & Sampling Seminar, October 21.
- "Skala: Accurate and scalable exchange-correlation with deep learning." National University of Singapore interdepartmental webinar, August 6.

2024

- "Diffusion Models." Nordic Probabilistic AI School, Copenhagen, June 17–21. Lecturer with Victor Garcia Satorras.
- "Diffusion Models." 7th Advanced Course on Data Science & Machine Learning, Tuscany, June 10–14. Plenary lecturer.

2023

- "Diffusion Models for Molecule Generation." Nordic Probabilistic AI School, Trondheim, June 12–16. Lecturer with Victor Garcia Satorras. [Materials](#)
- "Riemannian Diffusion Models." Learning on Graphs and Geometry / Starkly Speaking Reading Group, May 8. Invited talk with Milad Aghajohari.

2022

- "From Generative Flows to Diffusion Models." A*STAR Centre for Frontier AI Research seminar, November 25.
- "Convex Potential Flows, Diffusion Models, and Transition Paths." YES Workshop on Optimal Transport, Statistics, Machine Learning and Moving in Between, EURANDOM, Eindhoven, September 8. [Program](#)

2021

- "Optimal Transport and Probability Flows." Optimal Transport and Machine Learning Workshop, NeurIPS, December 13. Invited talk. [Workshop](#)

2020

- "Flows and Friends: Unifying Likelihood-Based Deep Probabilistic Models." Enriched Vision Applications Lab, National Chiao Tung University.

- "Improving Expressivity of Deep Probabilistic Models." Google Fellowship Summit.

2019

- "Roadmap to Expressive Flows." Google Cambridge Perception Team, October.
- "Go with the Flow: Recent Advances in Invertible Models." Element AI, May.

2018

- "Autoregressive Flows for Image Generation and Density Estimation." AI Summer School: Vision and Learning, National Tsing Hua University.
- "Autoregressive Flows for Image Generation and Density Estimation." Speech Processing and Machine Learning Lab, National Taiwan University.

SELECTED CONFERENCE PRESENTATIONS

- 2021: "A Variational Perspective on Diffusion-Based Generative Models and Score Matching." NeurIPS spotlight.
- 2021: "Convex Potential Flows." ICLR spotlight.
- 2020: "Solving ODE with Universal Flows: Approximation Theory for Flow-Based Models." ICLR workshop contributed talk.
- 2020: "Stochastic Neural Network with Kronecker Flow." AISTATS contributed talk.
- 2018: "Neural Autoregressive Flows." ICML oral.

TEACHING AND MENTORING

Internship supervision

- Benjamin Kurt Miller (2022), Microsoft Research Computational Catalysis internship: transition-state generation and reaction modeling.
- Bálint Máté (2023), Microsoft Research internship: architecture design exploration for learned density functionals.

Informal research mentoring

- Milad Aghajohari (2022): informally mentored while leading the Riemannian Diffusion Models project; co-developed the theory and codebase.
- Eeshan Dhekane (2019): mentored theory development for Hierarchical Importance Weighted Autoencoders.

Teaching

- Head Teaching Assistant, IFT6135 Representation Learning, Université de Montréal, 2019–2020.
- Teaching Assistant, IFT6135 Representation Learning, Université de Montréal, 2018.
- Co-developed assignments, projects, exams, tutorials, and lectures on generative modeling; taught tutorials on PyTorch, variational autoencoders, and normalizing flows.

WORKSHOP ORGANIZATION

- Organizer, [Advances in Programming Languages and Neurosymbolic Systems \(AIPANS\)](#), NeurIPS 2021.
- Lead Organizer, [Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models \(INNF+\)](#), ICML 2019–2021.

REVIEWING

- International Conference on Learning Representations (ICLR), 2021–2022.
- Neural Information Processing Systems (NeurIPS), 2019–2022 and 2025–2026.
- International Conference on Machine Learning (ICML), 2020–2023.
- ICML Workshop on Invertible Neural Networks and Normalizing Flows (INNF), 2019–2020.
- SN Partial Differential Equations and Applications, 2021.
- Association for the Advancement of Artificial Intelligence (AAAI), 2019.
- NeurIPS Workshop on Limited Labeled Data, 2017.

AWARDS

- [Google PhD Fellowship in Machine Learning](#), 2020–2022.
- FESP Bourse A de passage accéléré au doctorat (fast-track doctoral scholarship), Université de Montréal, 2018–2019.

LANGUAGES

Mandarin: native proficiency.

English: full professional proficiency.

French: limited working proficiency (B2 DELF).