

# ZHANG-WEI HONG

32 Vassar St, Cambridge, MA 02139, United States  
zwhong@mit.edu ♦ williamd4112.github.io

## EDUCATION

---

Massachusetts Institute of Technology

*start 2020 - 2024*

Ph.D. in Electrical Engineering and Computer Science,

Advised by *Prof. Pulkit Agrawal*

National Tsing Hua University

*start 2017 - end 2018*

Master in Computer Science,

Advised by *Prof. Chun-Yi Lee*

National Tsing Hua University

*start 2014 - end 2017*

Bachelor in Computer Science

## PUBLICATIONS

---

Jacob M. Desman, Zhang-Wei Hong, Moein Sabounchi, Ashwin S. Sawant, Jaskirat Gill, Ana C. Costa, Gagan Kumar, Rajeev Sharma, Arpeta Gupta, Paul McCarthy, Veena Nandwani, Doug Powell, Alexandra Carideo, Donnie Goodwin, Sanam Ahmed, Umesh Gidwani, Matthew A. Levin, Robin Varghese, Farzan Filsoufi, Robert Freeman, Avniel Shetreat-Klein, Alexander W. Charney, Ira Hofer, Lili Chan, David Reich, Patricia Kovatch, Roopa Kohli-Seth, Monica Kraft, Pulkit Agrawal, John A. Kellum, Girish N. Nadkarni, Ankit Sakhuja, **A distributional reinforcement learning model for optimal glucose control after cardiac surgery**, Nature, NPJ digital medicine

Guangtao Zeng\*, Maohao Shen\*, Delin Chen, Zhenting Qi, Subhro Das, Dan Gutfreund, David Cox, Gregory Wornell, Wei Lu, Zhang-Wei Hong\*, Chuang Gan, **Satori-SWE: Evolutionary Test-Time Scaling for Sample-Efficient Software Engineering**, Arxiv (\* denotes core contributors)

Sathwik Karnik\*, Zhang-Wei Hong\*, Nishant Abhangi\*, Yen-Chen Lin, Tsun-Hsuan Wang, Pulkit Agrawal, **Red Teaming Language-Conditioned Robot Models via Vision Language Models**, Accepted at NeurIPS Safe Generative AI Workshop 2024 (\* denotes co-first author)

Chi-Chang Lee\*, Zhang-Wei Hong\*, Pulkit Agrawal, **Going Beyond Heuristics by Imposing Policy Improvement as a Constraint**, Accepted at NeurIPS 2024 (\* denotes co-first author)

Srinath Mahankali, Zhang-Wei Hong, Ayush Sekhari, Alexander Rakhlin, Pulkit Agrawal, **Random Latent Exploration for Deep Reinforcement Learning**, Accepted at ICML 2024

Zhang-Wei Hong, Idan Shenfeld, Tsun-Hsuan Wang, Yung-Sung Chuang, Aldo Pareja, James R. Glass, Akash Srivastava, Pulkit Agrawal, **Curiosity-driven Red-teaming for Large Language Models**, Accepted at ICLR 2024

Srinath Mahankali\*, Chi-Chang Lee\*, Gabriel Margolis, Zhang-Wei Hong, Pulkit Agrawal, **Maximizing Quadruped Velocity by Minimizing Energy Models**, Accepted at ICRA 2024 (\* denotes co-first author)

Zhang-Wei Hong, Aviral Kumar, Sathwik Karnik, Abhishek Bhandwaldar, Akash Srivastava, Joni Pajarinen, Romain Laroche, Abhishek Gupta, and Pulkit Agrawal, **Beyond Uniform Sampling: Offline Reinforcement Learning with Imbalanced Datasets**, Accepted at *Conference on Neural Information Processing Systems (NeurIPS) 2023*

Idan Shenfeld, Zhang-Wei Hong, Aviv Tamar, and Pulkit Agrawal, **TGRL: Teacher-guided Reinforcement Learning for POMDP**, Accepted at *International Conference on Machine Learning (ICML) 2023*

Zechu Li, Tao Chen, Zhang-Wei Hong, Anurag Ajay, and Pulkit Agrawal, **Parallel Q-Learning: a Scheme for Time-efficient Reinforcement Learning**, Accepted at *International Conference on Machine Learning (ICML) 2023*

Zhang-Wei Hong, Pulkit Agrawal, Remi Tachet des Combes, and Romain Laroche, **Harnessing Mixed Offline Reinforcement Learning Datasets via Trajectory Reweighting**, Accepted at *International Conference on Learning Representation (ICLR) 2023*

Kwangjun Ahn, Zakaria Mhammedi, Horia Mania, Zhang-Wei Hong, and Ali Jadbabaie. **Model Predictive Control via On-Policy Imitation Learning**, Accepted as an oral presentation at *Learning for Decision Making and Control (L4DC) 2023*

Eric Chen\*, Zhang-Wei Hong\*, Joni Pajarinen, and Pulkit Agrawal. **Redeeming Intrinsic Rewards via Constrained Policy Optimization**, Accepted at *Conference on Neural Information Processing Systems (NeurIPS) 2022* (\* denotes co-first author)

Haokuan Luo, Albert Yue, Zhang-Wei Hong, Pulkit Agrawal. **Stubborn: A Strong Baseline for Indoor Object Navigation**, Accepted at *International Conference on Intelligent Robots and Systems (IROS) 2022*

Zhang-Wei Hong\*, Ge Yang\*, and Pulkit Agrawal. **Bilinear Value Networks for Multi-goal Reinforcement Learning**, Accepted at *International Conference on Learning Representation (ICLR) 2022* (\* denotes co-first author)

Zhang-Wei Hong, Tao Chen, Yen-Chen Lin, Joni Pajarinen, and Pulkit Agrawal. **Topological Experience Replay**, Accepted at *International Conference on Learning Representation (ICLR) 2022*

Chin-Jui Chang, Yu-Wei Chu, Chao-Hsien Ting, Hao-Kang Liu, Zhang-Wei Hong, and Chun-Yi Lee, **Reducing the Deployment-Time Inference Control Costs of Deep Reinforcement Learning Agents via an Asymmetric Architecture**, Accepted by *International Conference on Robotics and Automation (ICRA) 2021*

Zhang-Wei Hong, Prabhat Nagarajan, and Guilherme Maeda, **Periodic Intra-Ensemble Knowledge Distillation for Reinforcement Learning**, Accepted by *European Conference on Machine Learning (ECML) 2021* and *Deep Reinforcement Learning Workshop at Conference on Neural Information Processing Systems (NeurIPS) 2019*

Zhang-Wei Hong, Tsu-Jui Fu, Tzu-Yun Shann, Yi-Hsiang Chang, and Chun-Yi Lee. **Adversarial Active Exploration Strategy for Inverse Dynamics Model Learning**, Accepted as an oral paper by *Conference on Robot Learning (CoRL) 2019*

Zhang-Wei Hong, Tzu-Yun Shann, Shih-Yang Su, Yi-Hsiang Chang, Tsu-Jui Fu, and Chun-Yi Lee. **Diversity-driven Exploration Strategy for Deep Reinforcement Learning**, Accepted as a poster paper by *Conference on Neural Information Processing Systems (NeurIPS) 2018*

Zhang-Wei Hong, Chen Yu-Ming, Shih-Yang Su, Tzu-Yun Shann, Yi-Hsiang Chang, Hsuan-Kung Yang, Brian Hsi-Lin Ho, Chih-Chieh Tu, Yueh-Chuan Chang, Tsu-Ching Hsiao, Hsin-Wei Hsiao, Sih-Pin Lai, and Chun-Yi Lee **Virtual-to-Real: Learning to Control in Visual Semantic Segmentation**, Accepted as an oral paper by *International Joint Conferences on Artificial Intelligence (IJCAI) 2018*

Zhang-Wei Hong\*, Shih-Yang Su\*, Tzu-Yun Shann\*, Yi-Hsiang Chang, and Chun-Yi Lee. **Deep Policy Inference Q-Network for Multi-Agent Systems**, Accepted as an oral paper by *International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2018*

Yen-Chen Lin, Zhang-Wei Hong, Yuan-Hong Liao, Meng-Li Shih, Ming-Yu Liu, and Min Sun. **Tactics of adversarial attack on deep reinforcement learning agents**, Accepted as an oral paper by *International Joint Conferences on Artificial Intelligence (IJCAI) 2017*

## EXPERIENCE

---

<b>Research scientist</b> , MIT-IBM Research, Cambridge, MA, US	<i>2025 Jan. - present.</i>
<b>Research intern</b> , MIT-IBM Research, Cambridge, MA, US	<i>2023 Jun. - 2023 Sep.</i>
<b>Remote research intern</b> , Microsoft Research, Montreal, Canada	<i>2022 Jun. - 2022 Aug.</i>
<b>Graduate researcher / Graduate research assistant</b> , MIT, Cambridge	<i>2020 Sep. - Present</i>
<b>Full-time research assistant</b> , National Tsing Hua University, Taiwan	<i>2019 Oct. - 2020 Mar.</i>
<b>Research intern</b> , Preferred Networks, Japan	<i>2019 Jun. - 2019 Sep.</i>
<b>Engineering intern</b> , Appier, Taiwan	<i>2019 Feb. - 2019 Jun.</i>
<b>Visiting researcher</b> , Advised by <i>Prof. Jan Peters</i> , TU Darmstadt, Germany	<i>2018 Jul. - 2018 Sep.</i>
<b>Graduate research assistant</b> , National Tsing Hua University, Taiwan	<i>2016 Oct. - 2019 Jan.</i>
<b>Engineering intern</b> , Mediatek, Taiwan	<i>2016 Jul. - 2016 Sep.</i>
<b>Contract engineer</b> , Industrial Technology Research Institute, Taiwan	<i>2015 Oct. - 2015 Dec.</i>

## TEACHING

---

<b>6.484 - Computational Sensorimotor Learning</b> , MIT, U.S. Textbook drafting	<i>2022 Feb. - 2022 May.</i>
<b>6.S090 - Deep Learning for Control</b> , MIT, U.S. Lectures of off-policy reinforcement learning	<i>2021 Jan.</i>
<b>Nvidia deep learning institute</b> , Nvidia, Taiwan Hands-on image recognition	<i>2017 Jul. - 2017 Oct.</i>

## SERVICE

---

**International Conference on Learning Representation (ICLR)**, *Reviewer*  
**International Conference on Machine Learning (ICML)**, *Reviewer*  
**Conference on Robot Learning (CoRL)**, *Reviewer*  
**Conference on Neural Information Processing Systems (NeurIPS)**, *Reviewer*  
**International Conference on Intelligent Robots and Systems (IROS)**, *Reviewer*  
**Advanced Robotics Journal**, *Reviewer*  
**Goal-conditioned RL (GCRL) workshop**, NeurIPS, *Reviewer*  
**Foundational Models for Decision Making (FMDM) workshop**, NeurIPS, *Reviewer*  
**Deep RL workshop**, NeurIPS, *Program Committee*

## AWARDS AND SCHOLARSHIPS

---

<b>Qualcomm Innovation Fellowship</b> <sup>1</sup> <i>Qualcomm, US</i>	<i>2024</i>
<b>DAAD &amp; MOST Summer Institute Program Fellowship</b> <i>Ministry of Science and Technology and Deutscher Akademischer Austausch Dienst</i>	<i>2019</i>
<b>Nvidia Jetson Developer Challenge – World champion</b> <sup>23</sup> <i>Nvidia</i>	<i>2018</i>
<b>Nvidia Embedded Intelligent Robotics Challenge - 1st prize</b>	

---

<sup>1</sup><https://www.qualcomm.com/research/university-relations/innovation-fellowship/2024-north-america>

<sup>2</sup><https://challengerocket.com/nvidia/works/Sim-to-Real-Autonomous-Robotic-Control-adff14>

<sup>3</sup><https://insidebigdata.com/2018/04/10/winners-nvidiar-jetson-developer-challenge/>

**SKILLS**

---

**Programming Languages and Frameworks**

- C/C++/C#/Python/Java
- Message Passing Interface (MPI)/CUDA/OpenGL/Robot Operating System(ROS)/Ray
- Tensorflow/PyTorch/Chainer/Jax

**Languages**

- Mandarin (Chinese)
- English