

Mingdong Wu

Email: wmingd@pku.edu.cn

Homepage: aaronanima.github.io

Google Scholar: scholar.google.com/citations?user=MPfBNuIAAAAJ&hl=en

Research Interests: Architecting self-evolving robotic foundation models and large-scale distributed robot learning systems, through a robotic data-flywheel that autonomously collects real-world data to drive continuous self-improvement and autonomous skill acquisition.



EDUCATION

- Peking University** Beijing, China
• *Bachelor of Science - Turing Class of 2021, EECS; **Summa Cum Laude*** Sept. 2017 - Jun. 2021
Turing Advisors includes John Hopcroft, Manuel Blum, Silvio Micali, and Richard S. Sutton.
Teaching Assistant: Mathematics for Information Science (Prof. John Hopcroft).
- Peking University** Beijing, China
• *Turing PhD - Center on Frontiers of Computing Science; Assoc. Prof. Hao Dong* Sept. 2021 - Present

INTERNSHIP

- Beijing Institute for General Artificial Intelligence (BIGAI)** Beijing, China
• *Research Intern, Multi-agent Group* Jul. 2022 - Aug. 2023
 - **Research Focus:** Learning gradient fields for object rearrangement planning.
 - **Achievements:** Received BIGAI's Rising Star Award (1st Place) for outstanding contributions.
- Shanghai Zhiyuan Newcrek Technology Co., Ltd. (AgiBot)** Beijing and Shanghai, China
• *Research Intern, CTO-Office* Oct. 2023 - Feb. 2025
 - **Research Focus:** Real-world robotic reinforcement learning with tactile sensing.
 - **Achievements:** Received Agibot's Outstanding Contribution Award; Published two research papers in collaboration with PsiBot and XENSE Robotics; Pioneered the company's first industrial-level Proof of Concept (POC).
- Shanghai Zhiyuan Newcrek Technology Co., Ltd. (AgiBot)** Beijing and Shanghai, China
• *Research Intern, Directly Report to Jianlan Luo, Leading Scientist of AgiBot* Feb. 2025 - Sept. 2025
 - **Research Focus:** The company's top-priority strategic initiative; the details remain confidential.
 - **Achievements:** Primary project contributor.

PUBLICATIONS

(*) indicates equal contribution. **Mingdong Wu** is highlighted.

- **SimLauncher: Launching Sample-Efficient Real-world Robotic Reinforcement Learning via Simulation**
Pre-training:
Mingdong Wu*, Lehong Wu*, Yizhuo Wu*, Weiyao Huang, Hongwei Fan, Zheyuan Hu, Haoran Geng, Jinzhou Li, Jiahe Ying, Long Yang, Yuanpei Chen, and Hao Dong, *IROS 2025 (Oral)*.
- **UniTac2Pose: A Unified Approach Learned in Simulation for Generalizable Visuotactile In-hand Pose Estimation:**
Mingdong Wu*, Long Yang*, Jin Liu, Weiyao Huang, Lehong Wu, Zelin Chen, Daolin Ma, and Hao Dong, *CoRL 2025*.
- **DualGF: Example-driven Planning via Dual Gradient Fields:**
Mingdong Wu, Fangwei Zhong, Yulong Xia, Yizhou Wang, and Hao Dong, *ICANN 2025*.
- **Adaptive Visuo-Tactile Fusion with Predictive Force Attention for Dexterous Manipulation:**
Jinzhou Li*, Tianhao Wu*, Jiyao Zhang, Zeyuan Chen, Haotian Jin, **Mingdong Wu**, Yujun Shen, Yaodong Yang, and Hao Dong, *IROS 2025*.
- **Canonical representation and force-based pretraining of 3d tactile for dexterous visuo-tactile policy learning:**
Tianhao Wu, Jinzhou Li*, Jiyao Zhang*, **Mingdong Wu**, and Hao Dong, *ICRA 2025*.
- **Adamanip: Adaptive articulated object manipulation environments and policy learning:**
Yuanfei Wang, Xiaojie Zhang, Ruihai Wu, Yu Li, Yan Shen, **Mingdong Wu**, Zhaofeng He, Yizhou Wang, and Hao Dong, *ICLR 2025*.
- **Cordvip: Correspondence-based visuomotor policy for dexterous manipulation in real-world:**
Yankai Fu, Qiuxuan Feng, Ning Chen, Zichen Zhou, Mengzhen Liu, **Mingdong Wu**, Tianxing Chen, Shanyu Rong, Jiaming Liu, Hao Dong, and others, *RSS 2025*.
- **Boosting Universal LLM Reward Design through Heuristic Reward Observation Space Evolution:**
ZK Heng, Z Zhao, T Wu, Y Wang, **Mingdong Wu**, Y Wang, H Dong, *arXiv preprint arXiv:2504.07596, 2025*.
- **GFPack++: Improving 2D Irregular Packing by Learning Gradient Field with Attention:**
Tianyang Xue, Lin Lu, Yang Liu, **Mingdong Wu**, Hao Dong, Yanbin Zhang, Renmin Han, and Baoquan Chen, *ICCV 2025 (Highlight)*.
- **Omni6DPose: A Benchmark and Model for Universal 6D Object Pose Estimation and Tracking:**
Jiyao Zhang*, Weiyao Huang*, Bo Peng*, **Mingdong Wu**, Fei Hu, Zijian Chen, Bo Zhao, and Hao Dong, *ECCV 2024*.
- **Distilling Functional Rearrangement Priors from Large Models:**
Mingdong Wu*, Yiming Zeng*, Long Yang, Jiyao Zhang, Hao Ding, Hui Cheng, and Hao Dong, *RAL 2024*.

- **Mo-ddn: A coarse-to-fine attribute-based exploration agent for multi-object demand-driven navigation:**
Hongcheng Wang, Peiqi Liu, Wenzhe Cai, **Mingdong Wu**, Zhengyu Qian, and Hao Dong, *NeurIPS 2024*.
- **UniDexFPM: Universal dexterous functional pre-grasp manipulation via diffusion policy:**
Tianhao Wu, Yunchong Gan, **Mingdong Wu**, Jingbo Cheng, Yaodong Yang, Yixin Zhu, and Hao Dong, *arXiv preprint arXiv:2403.12421, 2024*.
- **SocialGFs: Learning social gradient fields for multi-agent reinforcement learning:**
Qing Long, Fangwei Zhong, **Mingdong Wu**, Yizhou Wang, and Song-Chun Zhu, *arXiv preprint arXiv:2405.01839, 2024*.
- **GenPose: Generative category-level object pose estimation via diffusion models:**
Mingdong Wu*, Jiyao Zhang*, and Hao Dong, *NeurIPS 2023*.
- **GraspGF: Learning Score-based Grasping Primitive for Human-assisting Dexterous Grasping:**
Mingdong Wu*, Tianhao Wu*, Jiyao Zhang, Yunchong Gan, and Hao Dong, *NeurIPS 2023*.
- **Find what you want: Learning demand-conditioned object attribute space for demand-driven navigation:**
Hongcheng Wang, Andy Guan Hong Chen, Xiaoqi Li, **Mingdong Wu**, and Hao Dong, *NeurIPS 2023*.
- **GFPack: Learning gradient fields for scalable and generalizable irregular packing:**
Mingdong Wu*, Tianyang Xue*, Lin Lu, Haoxuan Wang, Hao Dong, and Baoquan Chen, *SIGGRAPH Asia 2023*.
- **Score-PA: Score-based 3d part assembly:**
Junfeng Cheng, **Mingdong Wu**, Ruiyuan Zhang, Guanqi Zhan, Chao Wu, and Hao Dong, *BMVC 2023 (Oral)*.
- **GFPose: Learning 3d human pose prior with gradient fields:**
Hai Ci, **Mingdong Wu**, Wentao Zhu, Xiaoxuan Ma, Hao Dong, Fangwei Zhong, and Yizhou Wang, *CVPR 2023*.
- **Learning semantic-agnostic and spatial-aware representation for generalizable visual-audio navigation:**
Hongcheng Wang, Yuxuan Wang, Fangwei Zhong, **Mingdong Wu**, Jianwei Zhang, Yizhou Wang, and Hao Dong, *RAL 2023*.
- **TarGF: Learning target gradient field to rearrange objects without explicit goal specification:**
Mingdong Wu*, Fangwei Zhong*, Yulong Xia, and Hao Dong, *NeurIPS 2022*.

MANUSCRIPTS

(*) indicates equal contribution.

- **Boosting Universal LLM Reward Design through Heuristic Reward Observation Space Evolution:**
ZK Heng, Z Zhao, T Wu, Y Wang, **Mingdong Wu**, Y Wang, H Dong, *arXiv preprint arXiv:2504.07596, 2025*.
- **UniDexFPM: Universal dexterous functional pre-grasp manipulation via diffusion policy:**
Tianhao Wu, Yunchong Gan, **Mingdong Wu**, Jingbo Cheng, Yaodong Yang, Yixin Zhu, and Hao Dong, *arXiv preprint arXiv:2403.12421, 2024*.
- **SocialGFs: Learning social gradient fields for multi-agent reinforcement learning:**
Qing Long, Fangwei Zhong, **Mingdong Wu**, Yizhou Wang, and Song-Chun Zhu, *arXiv preprint arXiv:2405.01839, 2024*.

INVITED TALKS

- **6th Tech Sharing Forum, Agibot Robotics (Dec. 2024):**
Topic: Real-world Robotic Reinforcement Learning: Technology and Development Roadmap
Audience: CEO, CTO, CSO, and 200+ key technical staff
Previous Speakers: Including former founding partners.
- **Turing Student Research Forum (May 2023):**
Lightning Talk: Learning Target Gradient Fields for Object Rearrangement
Tutorial: Embracing Diffusion Models and Gradient Field Planners
- **3sigma Startup Roundtable Forum, 5Y Capital (Oct. 2022):**
Lightning Talk: Challenges for General AI in Robotics
Panelists: Senior Engineers from Google Brain, VPs from Xiaomi and JD.com
- **ZHIDX AI New Youth Talk (Nov. 2022):**
Live Talk: Learning Target Gradient Fields for Object Rearrangement
- **Beijing Institute for General Artificial Intelligence (Nov 2022):**
Long Talk: Learning Target Gradient Fields for Object Rearrangement
- **CMU Safe AI Laboratory (Oct. 2022):**
Long Talk: Learning Gradient Fields for Object Rearrangement

HONORS AND AWARDS

- **2025:** Turing Program Scholarship Award
- **2024:** Outstanding Contribution Award, Agibot
- **2023:** Best Tutorial Award & Best Poster Award, Turing Student Research Forum
- **2023:** BIGAI's Rising Star Award, 1st Place

PROFESSIONAL SERVICES

- **Reviewers:** T-RO, NeurIPS, ICLR, ICML, CVPR, RAL, CoRL, ICRA, IROS
- **Workshop Organizers:** 1st ICCV Workshop on "Category-Level Object Pose Estimation in the Wild", 2025