

# Xingdong Zuo

📍 Seongnam, South Korea    ✉️ [zuoxingdong@hotmail.com](mailto:zuoxingdong@hotmail.com)    🌐 [zuoxingdong.github.io](https://zuoxingdong.github.io)    📄 [zuoxingdong](#)  
📄 [zuoxingdong](#)    🎓 [Google Scholar](#)

---

## INTERESTS

Offline Reinforcement Learning for Recommender Systems and Auto-Bidding, Robot Learning with Vision-Language-Action Models, World Action Models

---

## EXPERIENCE

May 2021 – present    **Research Engineer**, Naver Corp – South Korea

Research and development of real-world reinforcement learning algorithms

- Developed auto-bidding for display ads using BC and IQL (offline RL), achieving significant improvement in ads efficiency; presentation at [Naver DAN 2025](#).
- Improved CVR and ROAS of NAVER Dynamic AD with advanced retrieval and RL models.
- Developed RL recommendation models for LINE Wallet, boosting long-term user engagement.
- Integrated offline RL into ranking models for Smart Message.
- Enhanced embedding and reranker models for Brand Chatbot and GAIA Shopping.
- Prototyped RLHF for user foundation models.
- Co-authored a NeurIPS 2023 publication and the HyperCLOVA X Technical Report; filed a Korean patent (2025).
- Built QueryLake ETL library and enhanced KPI reporting pipelines.

May 2017 – Nov 2017    **Master Thesis**, Bosch Center for Artificial Intelligence – Germany

Title: Sample-efficient model-based policy search via Bayesian Neural Networks Dynamics Models

- Achieved better sample efficiency: from ~10 to ~6 trials

---

## EDUCATION

2015 – 2019    **University of Freiburg**, M.Sc. in Computer Science – Freiburg, Germany  
Focus: Machine Learning and Reinforcement Learning

2010 – 2015    **Linnaeus University**, B.Sc. in Applied Mathematics – Sweden  
Minor: Computer Science

---

## SKILLS

**AI Stack:** Claude Code, PyTorch, Polars, Kubeflow, Docker, Trino SQL, Qdrant, Linux, Git

**Robotics:** LeRobot, Vision-Language-Action models (SmoIVLA), ACT, SO-101 / LeKiwi platforms

**Languages:** English (fluent), Chinese (native)

---

## PUBLICATIONS

- 2026 **Constrained Auto-Bidding via Generative Response Modeling**  
E Yang, X Zuo, K M Kim  
[arxiv.org/abs/2605.27811](https://arxiv.org/abs/2605.27811) [↗](#) (ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD))
- 2026 **Drift-VLA: Fast Vision-Language-Action Policies with One-Step Drifting**  
X Zuo  
[zuoxingdong.github.io/drift-vla](https://zuoxingdong.github.io/drift-vla) [↗](#) (Preprint)
- 2025 **Method and System to Provide Information Data with Minimal Infrastructure Resources**  
X Zuo, et al.  
[patents.google.com/patent/KR20250124713A/en](https://patents.google.com/patent/KR20250124713A/en) [↗](#) (Korean Patent Publication No. 10-2025-0124713)
- 2024 **HyperCLOVA X Technical Report**  
HyperCLOVA X Team (incl. X Zuo)  
[arxiv.org/abs/2404.01954](https://arxiv.org/abs/2404.01954) [↗](#) (arXiv preprint)
- 2023 **Direct preference-based policy optimization without reward modeling**  
G An, J Lee, X Zuo, N Kosaka, K M Kim, H O Song  
[arxiv.org/abs/2301.12842](https://arxiv.org/abs/2301.12842) [↗](#) (Advances in Neural Information Processing Systems 36)
- 2021 **Is Learning World Model Always Beneficial For Reinforcement Learning?**  
X Zuo  
[10.13140/RG.2.2.33768.80641](https://arxiv.org/abs/2110.13140) [↗](#) (Technical report)
- 2014 **Numerical simulation of Asano-Khrennikov-Ohya quantum-like decision making model**  
X Zuo  
(NeuroQuantology 12 (4))

---

## TALKS

- 2025 **Bid, Win, Profit: Revenue Optimization with a Generative-AI Ad-Auction Simulator – Seoul**  
Invited talk at DAN 25, NAVER’s annual tech conference ([recording](#) [↗](#)).

---

## PROJECTS

- 2026 – present **Mobile Manipulation with LeKiwi and PincOpen – Hugging Face**  
End-to-end mobile manipulation on LeKiwi (SO-101 arm, holonomic base); workflow published as a [Hugging Face blog](#) [↗](#).
- Fine-tuned SmolVLA for pick-and-place with data-scaling and vision-encoder ablations; the unfrozen encoder learned base navigation from only 10 of 130 episodes.
  - Integrated PincOpen gripper, upgraded servos, and fixed an arm-extension voltage-drop failure via P-gain tuning.
  - Published reproducible SmolVLA evaluation on LIBERO-Spatial ([blog](#) [↗](#)).
- 2025 – 2025 **LeRobot Worldwide Hackathon 2025 – Hugging Face**

- Led a 5-person team over 2 days to complete a language-instructed manipulation task using SmolVLA.
    - Fine-tuned SmolVLA for language-instructed color-ball pickup on real hardware
    - Led team coordination, task decomposition, and end-to-end workflow
  - 2025 – 2025 **LeRobot (contributor) – Hugging Face**  
Contributions to video encoding and CLI tools.
    - Authored a merged feature for batched video encoding in dataset recording, adding a `video_encoding_batch_size` option plus cleanup and exception handling for interrupted sessions ([PR #1390](#)).
  - 2025 – present [lerobokinson](#) & [lekiwi-tui](#) (author)  
GUI for professional PD gain tuning of Feetech/Dynamixel servos; terminal control for LeKiwi robots.
  - 2025 – present [rlbidder](#) (author)  
Modular reinforcement learning auto-bidding library for research and production.
  - 2018 – 2020 **OpenAI Gym (top #5 contributor, Nov. 2019)**  
42 merged PRs on wrappers and vectorized-environment infrastructure.
    - Implemented 12 (out of 14) different Wrappers classes
    - Refactored spaces and add new features: `flatdim`, `flatten`, `unflatten`
  - 2018 – 2020 **PyTorch (contributor)**
    - Implemented: `TanhTransform`, `DataLoader(drop_last)`, `convert Module params`, `ModuleList.insert`
    - Improved numerical stability in `SigmoidTransform`
  - 2018 – 2021 [lagom](#) (author)  
A PyTorch Library for Rapid Prototyping of Reinforcement Learning Algorithms (378★)
  - 2018 – 2020 [mazelab](#) (author)  
A customizable gym environment for maze/gridworld (271★, 13 citations).
  - 2017 – 2017 **[Value Iteration Networks](#) in PyTorch and TensorFlow**  
Widely used implementations for follow-up research (223★).
- Other open-source libraries (author)**  
[dm2gym](#) (87★, DeepMind Control Suite wrapper), [DeepPILCO](#) (54★, model-based RL with uncertainty), [recsys\\_metrics](#) (28★), [gym-recsys](#) (26★)