MINGHUI YE

my2035@hw.ac.uk Personal Website Github

EDUCATION

Heriot-Watt University Edinburgh, UK
Doctor of Philosophy 10/2023–10/2027 (expected)
Advisors: Prof. Ahmed H. Elsheikh, Dr. Nanzhe Wang
Research in two directions: (i) enhancing the data efficiency of model-based RL through uncertainty estimation and expressive policy classes (diffusion and flow policies); and (ii) applying model-based RL to engineering problems in subsurface flow control

Guangzhou University Guangzhou, China
Bachelor of Engineering in Robotics Engineering 09/2019–06/2023

SELECTED PUBLICATIONS

[1] **M. Ye**, A. H. Elsheikh. "DPMPC: Diffusion Policies as Expressive Priors for Model Predictive Control." *Under review, NeurIPS 2026.*

[2] **M. Ye**, A. H. Elsheikh. "Model-based reinforcement learning for active flow control." *Physics of Fluids*, 2025.

[3] **M. Ye**, J. Zhang. "Mobip: A lightweight model for driving perception using MobileNet." *Frontiers in Neurorobotics*, 2023.

EXPERIENCE

Doctoral Researcher, Heriot-Watt University, Edinburgh, UK 10/2023–present

- **DPMPC – Diffusion Policy Priors for Model Predictive Control:**
  - identified an underexplored bottleneck in high-dimensional MPC — the limited expressiveness of the Gaussian policy prior narrows the diversity of policy-guided proposals used to bootstrap the planner, which constrains exploration and undermines final performance
  - developed an efficient recipe for training a diffusion policy prior and incorporated it into the TDMPC framework
  - achieved substantial improvements over strong baselines across three challenging benchmarks (14 tasks on DMControl, MyoSuite and HumanoidBench), most markedly on humanoid control tasks with 61-dimensional action spaces
- **Model-Based RL for Active Flow Control:**
  - designed two control problems (state space, reward function, etc.) from scratch — synthetic-jet drag reduction on a cylinder and subsurface reservoir production maximization
  - adapted two model-based algorithms, planning (PETS) and Dyna-style RL (MBPO), to solve the flow control problems; improved the world model with a multi-step predictive loss, alleviating the compounding-error problem during autoregressive rollout
  - both MBRL algorithms outperformed the PPO baseline with 2–9$\times$ better sample efficiency and 2–7$\times$ faster convergence in wall-clock time; released code in the Gym-preCICE ecosystem

Visiting Researcher, Tongji University, Shanghai, China 09/2025–11/2025

- **Reduced-order world models for subsurface flow control:**

– used proper orthogonal decomposition (POD) to compress the high-dimensional flow state (pressure and velocity fields) into a compact latent representation, and learned the dynamics model directly in that latent space; the latent model was faster to roll out than the full-order model, enabling real-time planning and control

Visiting Researcher, NORCE Research, Bergen, Norway 07/2025–08/2025

- **Geosteering trajectory control under subsurface uncertainty:**
  – a drilling trajectory must be steered in real time to stay inside the reservoir, but the geology ahead of the drill bit is unknown and the agent sees only indirect, noisy gamma-ray logs — a partially observable control problem
  – adapted a world model to predict the gamma-ray response of the surrounding formation, attaining a probabilistic forecast model of the subsurface; coupled that forecast with uncertainty-aware model-based planning to reduce the risk of steering out of the reservoir

Research Assistant, Guangzhou University, Guangzhou, China 03/2022–09/2023

- **MobiP — Lightweight Multi-Task Driving Perception:**
  – built a lightweight multi-task perception network for autonomous driving that runs on embedded hardware rather than a dedicated GPU, widening where such algorithms can be deployed
  – designed an architecture with a single shared MobileNetV2 encoder and two-decoder heads — one detection head for traffic objects and one multi-class segmentation head covering drivable area and lane lines
  – with 7.5M parameters, MobiP matched the performance of YOLOP (the SOTA method) while running 45% faster, reaching 58 FPS on an NVIDIA Tesla V100 and 4.8 FPS on a Raspberry Pi 4B

HONORS & AWARDS

- James Watt Scholarship (fully-funded PhD), Heriot-Watt University 2023
- Second-Class Scholarship, Guangzhou University (top 5% of students), awarded twice 2020, 2023
- Third Prize, 12th Chinese National College Students Mathematics Competition 2020

INVITED TALKS

- "Model-Based Reinforcement Learning for Efficient Subsurface Flow Control," NORCE Research, Bergen, Norway 08/2025

TECHNICAL SKILLS

Programming: Python, C, MATLAB, Simulink, LaTeX, HTML/CSS
ML: PyTorch, JAX, OpenCV
Robotics: ABB IRB 1600, Arduino, Raspberry Pi, STM32, MuJoCo, Gazebo
Software: SLURM, Docker
Languages: Chinese (Native), English