Tong Zhou

Room 424, Ho Sin-Hang Engineering Bldg The Chinese University of Hong Kong Shatin, N.T., HKSAR, China
+86 13823524736 | zt1301112@gmail.com
Personal Webpage: https://www.timothijoe.com/projects

EDUCATION

The Chinese University of Hong Kong *Hong Kong, China*
Electronic Engineering – Ph.D 08/2017-Present

- Research Interests: Autonomous Driving, Decision Intelligence, Reinforcement Learning, Motion Planning,Monte Carlo Tree Search
- Relevant Coursework: Advance Theory of Probability, Pattern Recognition, Advanced Robotics, Matrix Analysis and Computation, Advanced Topics in Game Theory, Optimization Theory.

Harbin Institute of Technology, China (HIT) *Heilongjinag, China*
Measurement, Control Technique and Instruments - Bachelor of Engineering 08/2013-05/2017

- GPA: 90.99/100
- Rank: 3/114
- Relevant Coursework: Advanced Mathematics, Physical Electronics Control
- Honors:
  - National Scholarship for Academic Year 2014-2015
  - The Third-class of People's scholarship for Academic Year 2013-2014
  - The Second-class of People's scholarship for Academic Years 2014-2015 & 2015-2016

AWARDS

- T.J. Tran Best Paper in Robotics Award in 2021 International Conference on Robotics and Biomimetics
- Student Creativity Awards in 2019, issued by The Chinese University of Hong Kong – First Runner-up
- Obtained invention patent certificate for Airport Baggage Cart Automatic Recycling System in 2019
- NI LabVIEW Certified Associate Developer (CLAD) from 05/2016 to 05/2018
- 2014 China Undergraduate Mathematical Contest in Modeling in Heilongjiang Province - First Prize
- Bionic Robot Goal Seeking Contest of "HIT Cup" 16th National Robot Championship Contest and 'Bosiweilong Cup 5th International Humanoid Robot Olympic Games – First Prize

INTERNSHIP

Beijing SenseTime Technology Co.Ltd 35 hours/week 05/2021-07/2023
Contributed to the advancement of several projects within OpenDILab at Sensetime and Shanghai AI Lab.

- DI-drive: a specialized platform developed within OpenDILab, meticulously tailored for addressing driving-related tasks.
  - Main development tasks: revolved around the reproduction and implementation of algorithms, with a particular focus on leveraging the capabilities of the Carla and MetaDrive simulators.
- DI-engine: an official open-source decision intelligence platform, developed on PyTorch and JAX, tailored for reinforcement learning tasks in diverse environments.
  - Responsibility: replicated and implemented exploration-related reinforcement learning algorithms, such as Intrinsic Curiosity Module and Soft Q-Learning.
- LightZero: a formally recognized, lightweight, and efficient open-source algorithm library unifying Monte Carlo Tree Search with reinforcement learning principles.
  - Responsibility: developed MuZero, Sampled MuZero, and Stochastic MuZero algorithms within the LightZero library.

Neusoft Corporation Intern in Training Department 35 hours/week 09/2016-10/2016

- Collaborated in a 5-member team to compile 32 references on Google's website. Demonstrated proficiency in the operation and implementation of elevator monitoring systems. Acquired expertise in microcontroller programming using *Keil* and policy circuit diagram creation using *Proteus*.
- Proficiently utilized simulation techniques and leveraged *C programming* to successfully implement core functionalities of elevator monitoring, encompassing 'door operation', 'floor summoning', and 'emergency processing'."
- Utilized *Proteus* for circuit diagram design and optimization. Successfully executed the elevator monitoring system project.

RESEARCH EXPEROENCE

All Projects Details can be found in the link: https://www.timothijoe.com/projects

Minimally-Intrusive Navigation in Densely-Populated Pedestrian Flow 08/2021-Present

- Introduced the concept of 'Minimal Intrusive Navigation,' which involves assessing disturbances in pedestrian flow at both macroscopic and microscopic levels. Macroscopically, disturbances are quantified by factors like density and direction. Microscopically, it entails evaluating the discrepancy in future trajectory distributions for individual

pedestrians in the presence or absence of the robot, indicating the additional displacement required to avoid the robot's path.

- Pioneered a triangle-based sampling strategy, effectively subdividing the spatial surroundings of the robot into topologically defined 'inserting areas' tailored for precise sampling utilizing motion primitive methods.
- Engineered a comprehensive navigation workflow, accounting for intrusiveness at both macro and micro-level modeling. Implemented the integration of pedestrians' eye gaze data to assess their likelihood of perceiving the robot, enabling precise adjustments in micro-level pedestrian modeling.

Expert-Guided Motion-Encoding Tree Search in Autonomous Driving 2023/09-Present

- Integrated motion primitive methods into Monte Carlo Tree Search (MCTS) framework, effectively mitigating search complexity in extended-duration MCTS tasks by reducing tree depth
- Applied diverse expert policies to uphold a multi-modal policy distribution, ensuring the availability of multiple potential policies for selection even under identical observations.
- Implemented MCTS search with expert policy integration and skill-based execution, optimizing the application of expert strategies to maintain diversity in policy output and circumvent local optima.

Task Agnostic and Ego-Centric Reinforcement Learning in Autonomous Driving 08/2021-Present

- Engineered a task-agnostic and ego-centric motion skill library, demonstrating versatility in encompassing a wide array of motion skills, and facilitating cross-task reusability.
- Transformed motion skills into a latent skill space and tailored the RL algorithm to navigate within this domain, resulting in heightened proficiency and significant learning impact.
- Evidenced our approach's proficiency in achieving efficient and effective autonomous driving learning across three challenging dense-traffic scenarios.

LightZero 10/2022-07/2023

Contributed to the development of LightZero, an open-source algorithm toolkit integrating Monte Carlo Tree Search and Deep Reinforcement Learning.

Responsibilities:

- Reproduced the MuZero algorithm, an innovative method that extends the applicability of techniques akin to Alphago, enabling tree search in environments with unknown transition dynamics.
- Implemented the Sampled MuZero method, an extension of MuZero, to facilitate learning in domains with arbitrarily complex action spaces through strategic planning over sampled actions.
- Implemented the Stochastic MuZero method, enabling comprehensive incorporation of the stochastic nature of the environment in tree search processes.

DI-drive 05/2021-09/2022

Participated in the development of an open-source Decision Intelligence Platform for Autonomous Driving simulation.

Responsibilities:

- Integrated the Model Predictive Control algorithm into the CARLA simulator, seamlessly combining trajectory tracking and path following modes.
- Implemented macro-level functionalities within the MetaDrive Simulator's highway environment, employing the DQN algorithm to formulate lane-changing strategies.
- Replicated the Disturbance-based Reward Extrapolation (D-REX) algorithm in the Drive environment, enabling the vehicle to acquire effective strategies and surpass provided demonstrations, even in the presence of suboptimal inputs.

Online State-Time Trajectory Planning in Highly Dynamic Crowd Environments 05/2020-04/2021

- Proposed a gradient-based planner in the state-time domain, enabling real-time trajectory generation in dynamically demanding environments.
- Introduced timed ESDT arguments, facilitating distance and gradient inquiries using state-time keys for motion trajectory optimization. Formulated smooth prior and obstacle likelihood functions compatible with the state-time domain, transforming the trajectory planning problem into a mapping challenge, subsequently addressed via numerical optimization techniques.
- Verified the proposed method using simulated and benchmark datasets, with experimental results showcasing its superior performance over state-of-the-art approaches, thus underscoring its substantial advantages in contrast to conventional methodologies.

Human Trajectory Prediction 05/2019-05/2020

- Acquired and assessed the ETH and UCY pedestrian datasets, and gathered densely populated pedestrian data from Hong Kong's metro stations and plazas employing calibrated lidar-camera systems..
- Derived from data analysis, identified a tendency where individuals allocate greater attention to pedestrians in front, while exhibiting reduced attention towards those behind. Introduced a novel attention-based social pooling method in response.
- Developed encoders capable of forecasting pedestrians' future trajectories. Through comparison with existing LSTM and GAN techniques, demonstrated notable enhancements in both average displacement error and final displacement error.

Airport Baggage Cart Automatic Recycling System 05/2018-05/2019

- Mastered the intricacies of airport self-collection baggage systems, precisely determining the position and orientation of the target baggage trolley. Employed multi-sensor fusion techniques, including wheel odometer, 3D dimensional lidar, camera, and ultrasonic sensor, to achieve comprehensive trolley detection.
- Conducted thorough assessment of trolley fine attitude estimation, resulting in improved maneuvering capabilities. Implemented proximity vision waiting methodology to enable the robot to derive the optimal proximity action trajectory, effectively guiding the luggage trolley to its destination.
- Demonstrated significant cost savings in labor, particularly in developed regions grappling with pronounced aging demographics or elevated labor expenses.

An Autonomous Eye-in-Hand robotic System for Elevator Button Operation 10/2017-04/2018

- Conducted extensive literature review on elevator button operation systems, leading to the conceptualization of an autonomous robot system with Eye-in-Hand architecture to address the button operation challenge.
- Leveraged PyTorch to develop advanced deep neural networks, applied in button detection and character recognition tasks to furnish precise perceptual input. Additionally, introduced a button pose estimation algorithm, achieving accurate pose estimation through a minimal uncertainty model fitting process.
- Formulated a coarse-to-precise control framework to guide the manipulator in executing button operation tasks. This innovative approach was successfully published in the IEEE Transactions on Instrumentation and Measurement in 2020.

ACTIVITIES

- Multivariate Calculus For Engineers (ENGG 1130E) 09/2020-04/2021
  Tutor 10 hours/week
- Intelligent Interactive Robot Practice (ELEG 4701) 09/2019-08/2020
  Tutor 10 hours/week
- Introduction to Embedded System (ELEG 2401) 09/2018-08/2019
  Tutor 10 hours/week
- Engineering Physics: Electromagnetics, Optics and Modern Physics (ENGG 1301) 09/2017-08/2018
  Tutor 10 hours/week
- HIT Virtual Instrument Technology Association – Training Department 09/2015-08/2016
  Director 5 hours/week
- Student Union from EE and Automation College – Learning Department 09/2013-08/2014
  Staff 4 hours/week

PUBLICATIONS

Bibliography (Name of Tong Zhou is bold)

- **T Zhou**, E Lyu, J Wang, G Cen, Z Zha, S Qi, MQH Meng. Towards High Efficient Long-horizon Planning with Expert-guided Motion Encoding Tree Search. (arXiv preprint arXiv: 2309.15079) - Published in 2023
- **T Zhou**, L Wang, R Chen, W Wang, Y Liu. Accelerating Reinforcement Learning for Autonomous Driving using Task-Agnostic and Ego-Centric Motion Skills (2023 IEEE/RSJ International Conference on Intelligent Robots and Systems) - Published in 2023
- Y. Niu, Y. Pu, Z. Yang, X. Li, **T. Zhou**, J. Ren, S. Hu, H. Li, and Y. Liu. Lightzero: A unified benchmark for monte carlo tree search in general sequential decision scenarios (Conference on Neural Information Processing Systems) -Published in 2023
- D Zhu, **T Zhou**, J Lin, Y Fang, MQH Meng. Online State-Time Trajectory Planning Using Timed-ESDF in Highly Dynamic Environments. (2022 International Conference on Robotics and Automation) - Published in 2022
- J Lin, **T Zhou**, D Zhu, J Liu, MQH Meng. Search-Based Online Trajectory Planning for Car-Like Robots in High Dynamic Environments. (2021 IEEE International Conference on Robotics and Automation) - Published in 2021
- **T Zhou**, S Qi, E Lyu, G Cen, J Wang, MQH Meng. Towards Minimally-Intrusive Navigation in Densely-Populated Pedestrian Flow. (2021 IEEE International Conference on Robotics and Biomimetics) - Published in 2021
- D Zhu, Z Min, **T Zhou**, T Li, MQH Meng. An Autonomous Eye-in-Hand Robotic System for Elevator Button Operation Based on Deep Recognition Network. (IEEE Transactions on Instrumentation and Measurement 70, 1-13) - Published in 2020
- D Zhu, T Li, D Ho, **T Zhou**, MQH Meng. A novel OCR-RCNN for elevator button recognition. (2018 IEEE/RSJ International Conference on Intelligent Robots and Systems) - Published in 2018

ADDITIONAL INFORMATION

Programming C++(Expertise), Python (Expertise), Pytorch (Expertise) , ROS(expertise), MATLAB (Experience), LabVIEW(Experience), Lua(Experience), CAD(Experience)

Language Chinese (native)
English (fluent), IELTS Overall 6.5 (Reading: 8.0)

Interests Tennis (4 years), Reading (especially in history and Philosophy), Swimming (3 years)