Research Statement of Zihan Yang

Introduction. Motivated by the limitations of how current robots operate, I aim to develop generalizable solutions that can adapt to various tasks and environments. I believe that endowing robots with generalizable decision-making capabilities is key to unlocking their transformative impact on society.

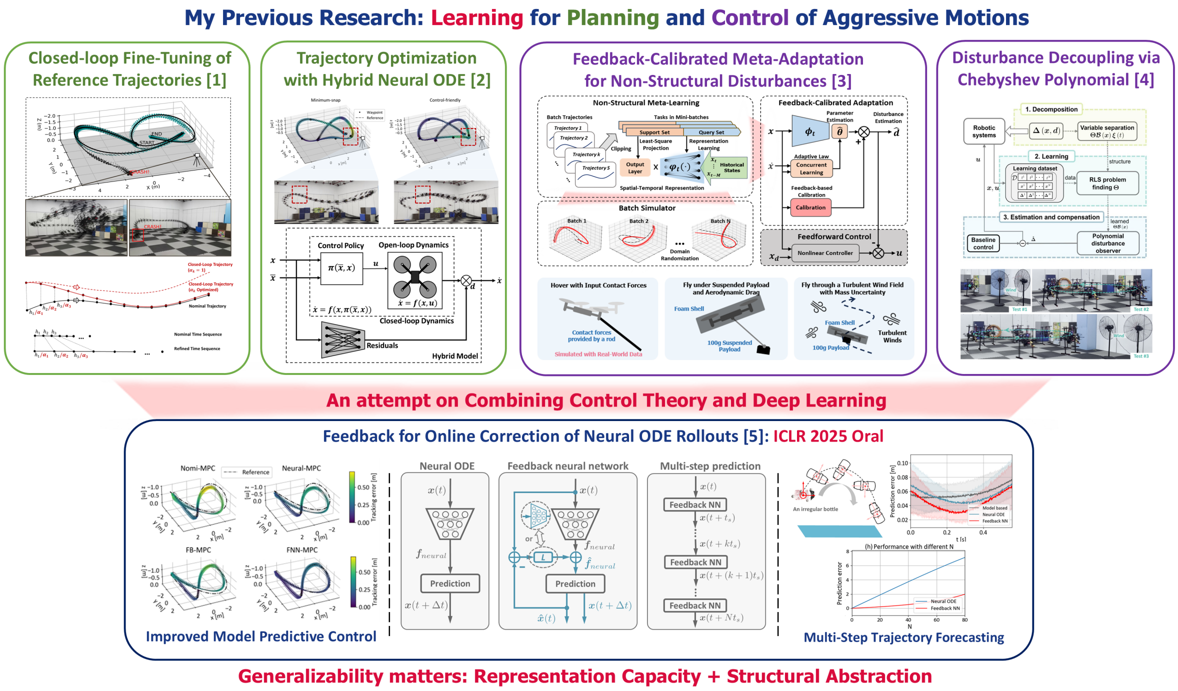

Figure 1: Overview of my previous research on learning-based planning and control, with extensions toward applying control theory to the design of deep learning models.

Previous Research. My research has focused on learning-based methods for taming residual physics and disturbances in robotic dynamics. I have developed several frameworks from both planning and control perspectives. On the planning side, my work involves fine-tuning reference trajectories [1] to minimize control errors (Best Student Paper Award in ICCA 2024) and leveraging neural ordinary differential equations (neural ODEs) to model and mitigate residual effects along optimized trajectories [2]. For control, I have explored online learning of disturbances using meta-learned representations that enable real-time adaptation to novel disturbances [3]. Additionally, I have investigated the decoupling of disturbances into state-dependent and time-varying components using Chebyshev polynomials [4]. Through these efforts, I observed that model generalizability stems from both **representation capacity** and **structural abstraction**, with the latter benefiting significantly from principles in control theory. Building on this insight, I proposed a feedback-augmented neural ODE framework that learns latent dynamics while incorporating online state feedback to calibrate model rollouts, resulting in more generalizable online predictions [5] (Oral in ICLR 2025). Despite promising results, these models are not yet directly applicable to more complex tasks, such as locomotion of articulated robots, long-horizon planning, or decision-making in manipulation. This motivates my ongoing efforts towards solving more general robotic tasks.

Current Interests. Generative models, such as diffusion-based [6], [7] and flow-based architectures [8], have demonstrated remarkable compositional generalization. When conditioned on specific tasks or environments, these models exhibit strong adaptability in robotic applications

[9], [10], [11]. I am developing methods to integrate generative models into model predictive control (MPC) frameworks, aiming to enhance the control robustness of quadrupedal robots across diverse terrains and payloads. This is part of my broader vision: developing generalizable models and policies that scale to general robotic applications.

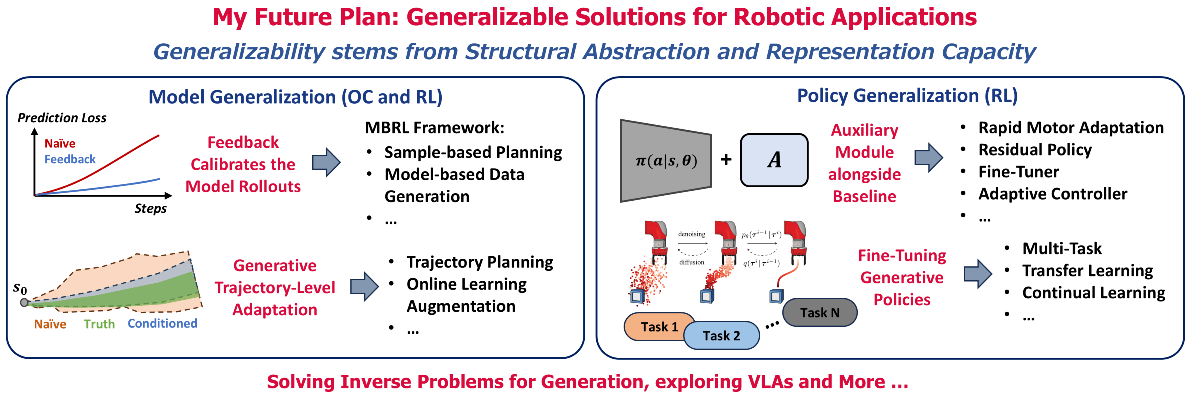

Figure 2: Research roadmap illustrating my future work on generalizable modeling in robotics, with extensions toward control-guided generative models and advanced embodied intelligence through vision-language-action frameworks.

Future Research. Optimal control (OC) and reinforcement learning (RL) are two promising approaches to solve control, planning, and higher-level decision-making problems in robotics [12]. Enhancing model generalization is crucial across both OC and RL frameworks, while policy generalization is uniquely critical in RL settings. My research tackles this challenge by focusing on two complementary themes: (i) leveraging structural insights to guide model and policy generalization, and (ii) incorporating expressive models, such as generative models, to enhance adaptability across diverse robotic tasks.

• **Model.** Models are central to robotics, providing key representations for planning and control. Incorporating structural feedback from control theory improves model generalization [5], especially in model-based RL [13], [14], by calibrating rollouts for more reliable planning. Using high-capacity generative models further boosts predictive performance [9], [11], [15]. Combining these with online learning can be a solution for better trajectory-level adaptation and generalization.

• **Policy.** Building adaptable policies for diverse tasks is a core challenge in robotics. This can be addressed by structurally enhancing baseline policies with modules like environment encoders [16], residual policy [17] or fine-tuning components [18], [19], and adaptive mechanisms [20], [21]. Generative models can further augment policy generalization [22], while task-specific fine-tuning enables multi-task and continual learning beyond imitation [23].

• **Extensions.** Orthogonal to the above two directions, I am also interested in fine-tuning generative models using control and optimization. The guidance of diffusion-like generative models can be seen as an inverse problem [24], [25], where a goal condition or posterior distribution is desired to be achieved. This opens up the new possibility of integrating generative models with control and optimization. In addition, exploring the integration of more advanced embodied foundation models, such as Vision-Language-Action Models [26], represents a promising direction. These models combine large-scale perception, reasoning, and decision-making capabilities, potentially enabling robots to perform a broader range of tasks with greater autonomy.

References

[1] **Z. Yang**, J. Jia, Y. Liu, K. Guo, X. Yu, and L. Guo, "TRACE: Trajectory refinement with control error enables safe and accurate maneuvers," in *IEEE 18th International Conference on Control and Automation*, **Best Student Paper Award**, 2024.

[2] K. Guo*, **Z. Yang***, J. Jia, Y. Liu, and X. Yu, "Optimizing control-friendly trajectories with self-supervised residual learning," *Under Review*, 2024.

[3] **Z. Yang**, J. Jia, M. Wang, Y. Liu, K. Guo, and X. Yu, "Towards non-structural disturbance prediction: Feedback-calibrated meta-adaptation," *Under Review*, 2025.

[4] J. Jia*, M. Wang*, **Z. Yang**, B. Yang, Y. Liu, K. Guo, and X. Yu, "Learning-based observer for coupled disturbance," 2025.

[5] J. Jia*, **Z. Yang***, M. Wang, K. Guo, J. Yang, X. Yu, and L. Guo, "Feedback favors the generalization of neural ODEs," in *The Thirteenth International Conference on Learning Representations*, **Oral Presentation (acceptance rate: 1.8%)**, 2025.

[6] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," 2020.

[7] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," 2021.

[8] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le, "Flow matching for generative modeling," 2023.

[9] M. Janner, Y. Du, J. B. Tenenbaum, and S. Levine, "Planning with diffusion for flexible behavior synthesis," in *Proceedings of the 39th International Conference on Machine Learning*, 2022.

[10] A. Ajay, Y. Du, A. Gupta, J. Tenenbaum, T. Jaakkola, and P. Agrawal, "Is conditional generative modeling all you need for decision-making?" In *The Eleventh International Conference on Learning Representations*, 2023.

[11] G. Zhou, S. Swaminathan, R. V. Raju, J. S. Guntupalli, W. Lehrach, J. Ortiz, A. Dedieu, M. Lázaro-Gredilla, and K. Murphy, "Diffusion model predictive control," 2024.

[12] Y. Song, A. Romero, M. Müller, V. Koltun, and D. Scaramuzza, "Reaching the limit in autonomous racing: Optimal control versus reinforcement learning," *Science Robotics*, vol. 8, no. 82, eadg1462, 2023.

[13] A. Nagabandi, I. Clavera, S. Liu, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn, "Learning to adapt in dynamic, Real-world environments through meta-reinforcement learning," in *The Seventh International Conference on Learning Representations*, 2019.

[14] N. A. Hansen, H. Su, and X. Wang, "Temporal difference learning for model predictive control," in *Proceedings of the 39th International Conference on Machine Learning*, 2022.

[15] J. Carvalho, A. T. Le, M. Baierl, D. Koert, and J. Peters, "Motion planning diffusion: Learning and planning of robot motions with diffusion models," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023.

[16] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "RMA: Rapid motor adaptation for legged robots," in *Proceedings of Robotics: Science and Systems*, 2021.

[17] X. Yuan, T. Mu, S. Tao, Y. Fang, M. Zhang, and H. Su, “Policy decorator: Model-agnostic online refinement for large policy model,” in *The Thirteenth International Conference on Learning Representations*, 2025.

[18] M. S. Mark, T. Gao, G. G. Sampaio, M. K. Srirama, A. Sharma, C. Finn, and A. Kumar, “Policy agnostic RL: Offline RL and online RL fine-tuning of any class and backbone,” 2024.

[19] M. Nakamoto, S. Zhai, A. Singh, M. Sobol Mark, Y. Ma, C. Finn, A. Kumar, and S. Levine, “Cal-QL: Calibrated offline RL pre-training for efficient online fine-tuning,” in *Advances in Neural Information Processing Systems*, 2023, pp. 62 244–62 269.

[20] S. Lyu, X. Lang, H. Zhao, H. Zhang, P. Ding, and D. Wang, “RL2AC: Reinforcement learning-based rapid online adaptive control for legged robot robust locomotion,” in *Proceedings of Robotics: Science and Systems*, 2024.

[21] C. S. Brito and D. C. McNamee, “World models as reference trajectories for rapid motor adaptation,” in *7th Robot Learning Workshop: Towards Robots with Human-Level Abilities, The Thirteenth International Conference on Learning Representations*, 2025.

[22] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. C. Burchfiel, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” in *Proceedings of Robotics: Science and Systems*, 2023.

[23] A. Z. Ren, J. Lidard, L. L. Ankile, A. Simeonov, P. Agrawal, A. Majumdar, B. Burchfiel, H. Dai, and M. Simchowitz, “Diffusion policy policy optimization,” in *The Thirteenth International Conference on Learning Representations*, 2025.

[24] C. Domingo-Enrich, M. Drozdzal, B. Karrer, and R. T. Q. Chen, “Adjoint matching: Fine-tuning flow and diffusion generative models with memoryless stochastic optimal control,” in *The Thirteenth International Conference on Learning Representations*, 2025.

[25] L. Wang, C. Cheng, Y. Liao, Y. Qu, and G. Liu, “Training free guided flow matching with optimal control,” in *The Thirteenth International Conference on Learning Representations*, 2025.

[26] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. P. Foster, P. R. Sanketi, Q. Vuong, T. Kollar, B. Burchfiel, R. Tedrake, D. Sadigh, S. Levine, P. Liang, and C. Finn, “OpenVLA: An open-source vision-language-action model,” in *8th Annual Conference on Robot Learning*, 2024.