

Motivation

Can we develop an end-to-end method that directly learns legible motion from offline demonstrations?



Legible Robot Motion [1] **Learning from Demonstrations** [2]

- Intent expressive
- Improved safety
- Faster task completion
- Safe offline training
- Non-expert programming
- Highly scalable

Experiment

Training Dataset

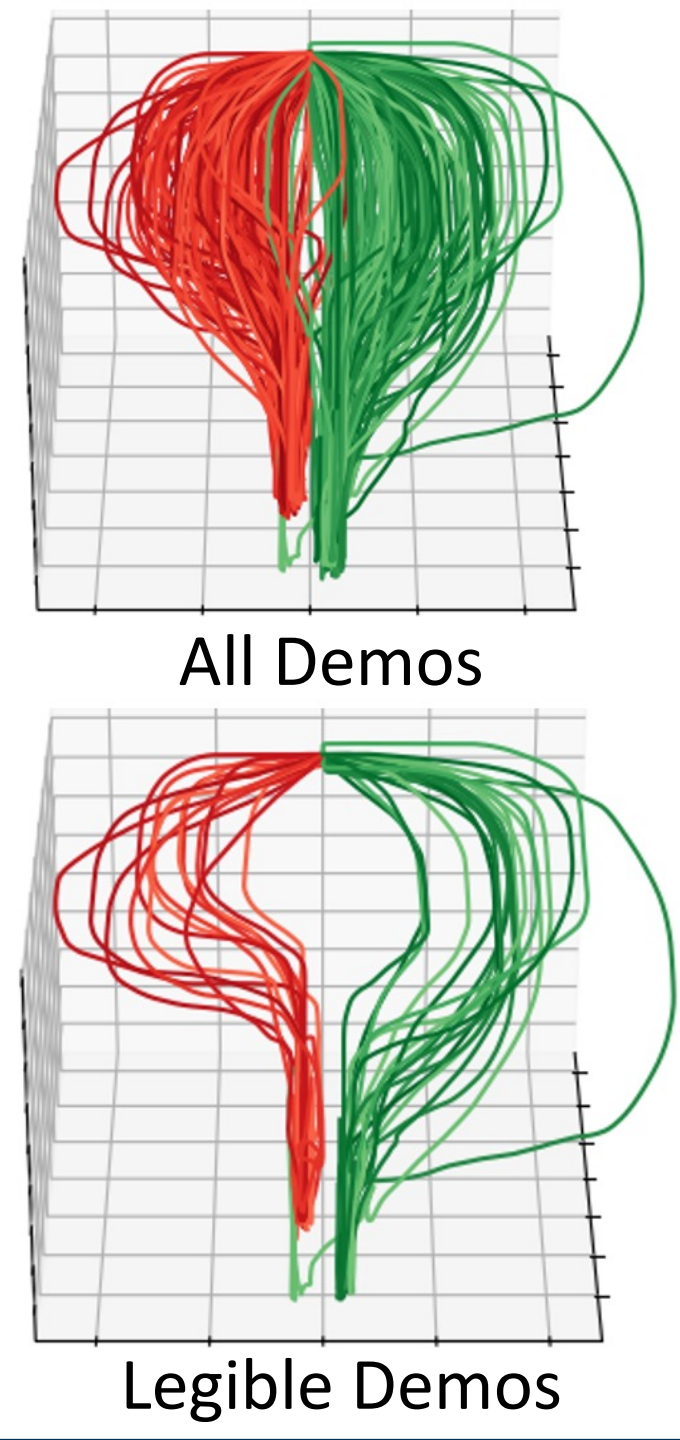
- 280 multi-modal demonstrations of successful block lifts

Generative Algorithms Evaluated

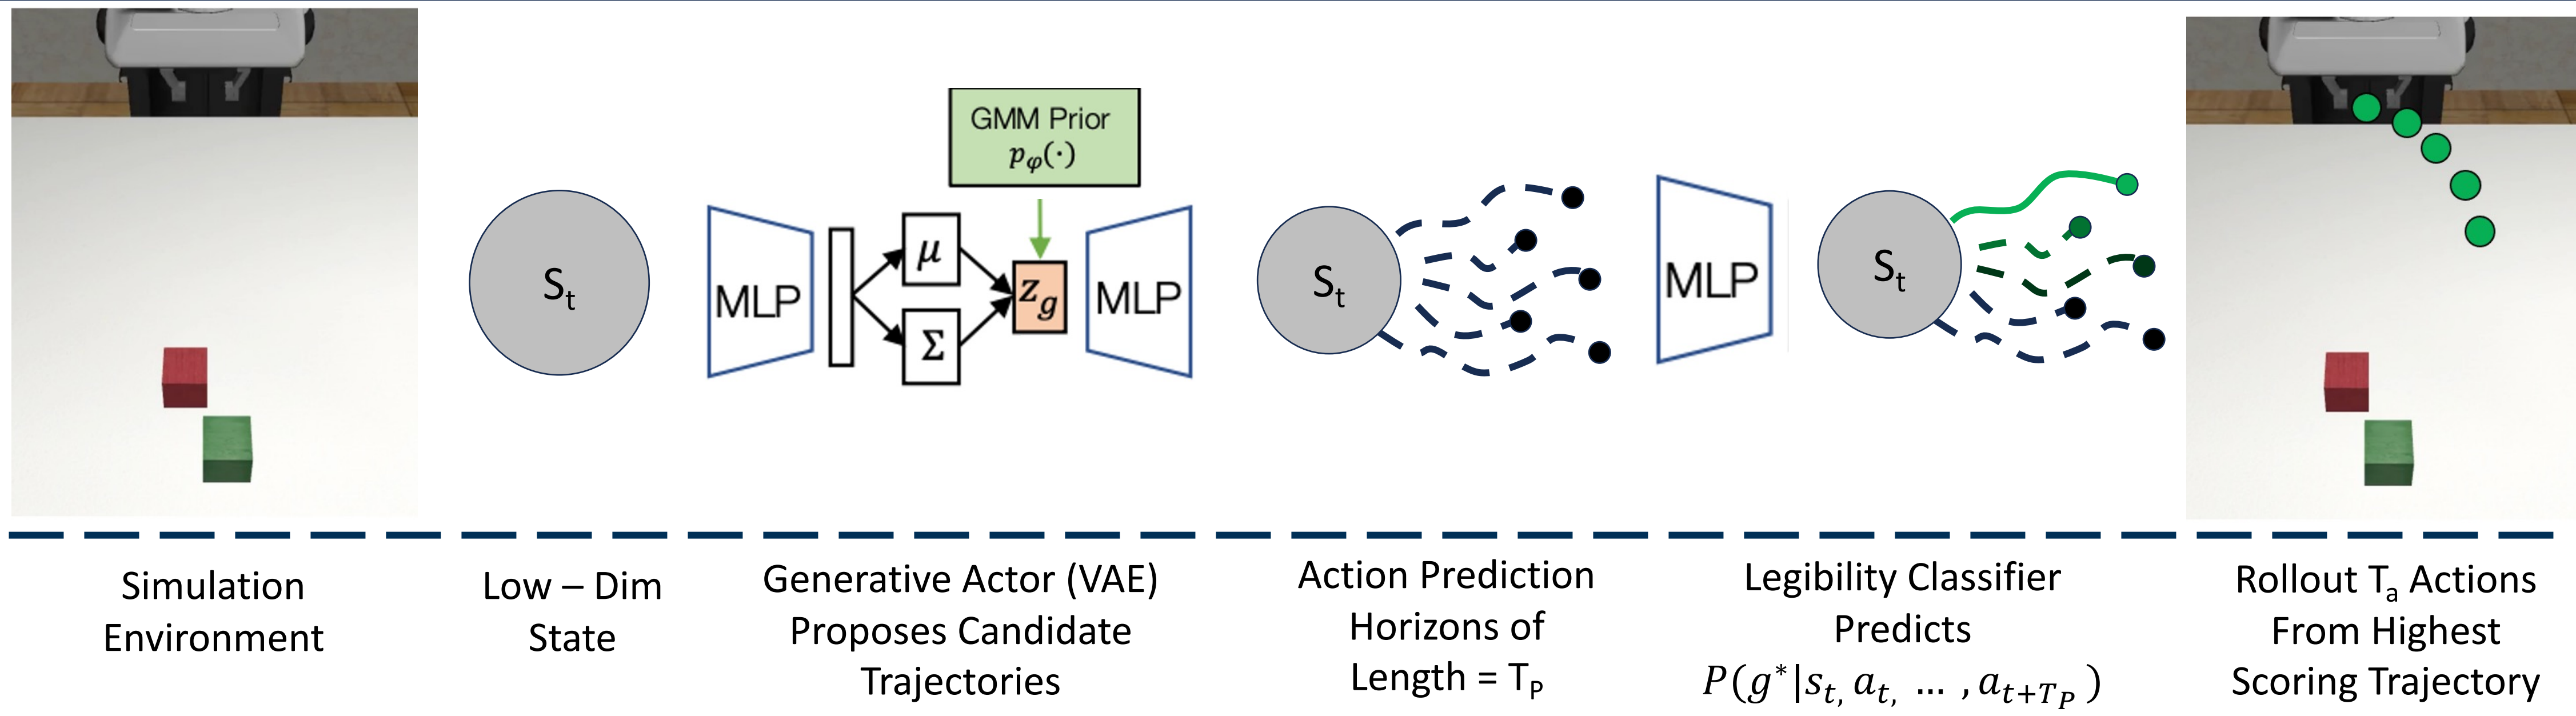
- Variational Auto-Encoder (VAE)
- Goal Conditioned VAE (G-VAE)
- *Generative Legible Motion Model (GLMM)*

Independent Variables

- Success Rate
- Legibility

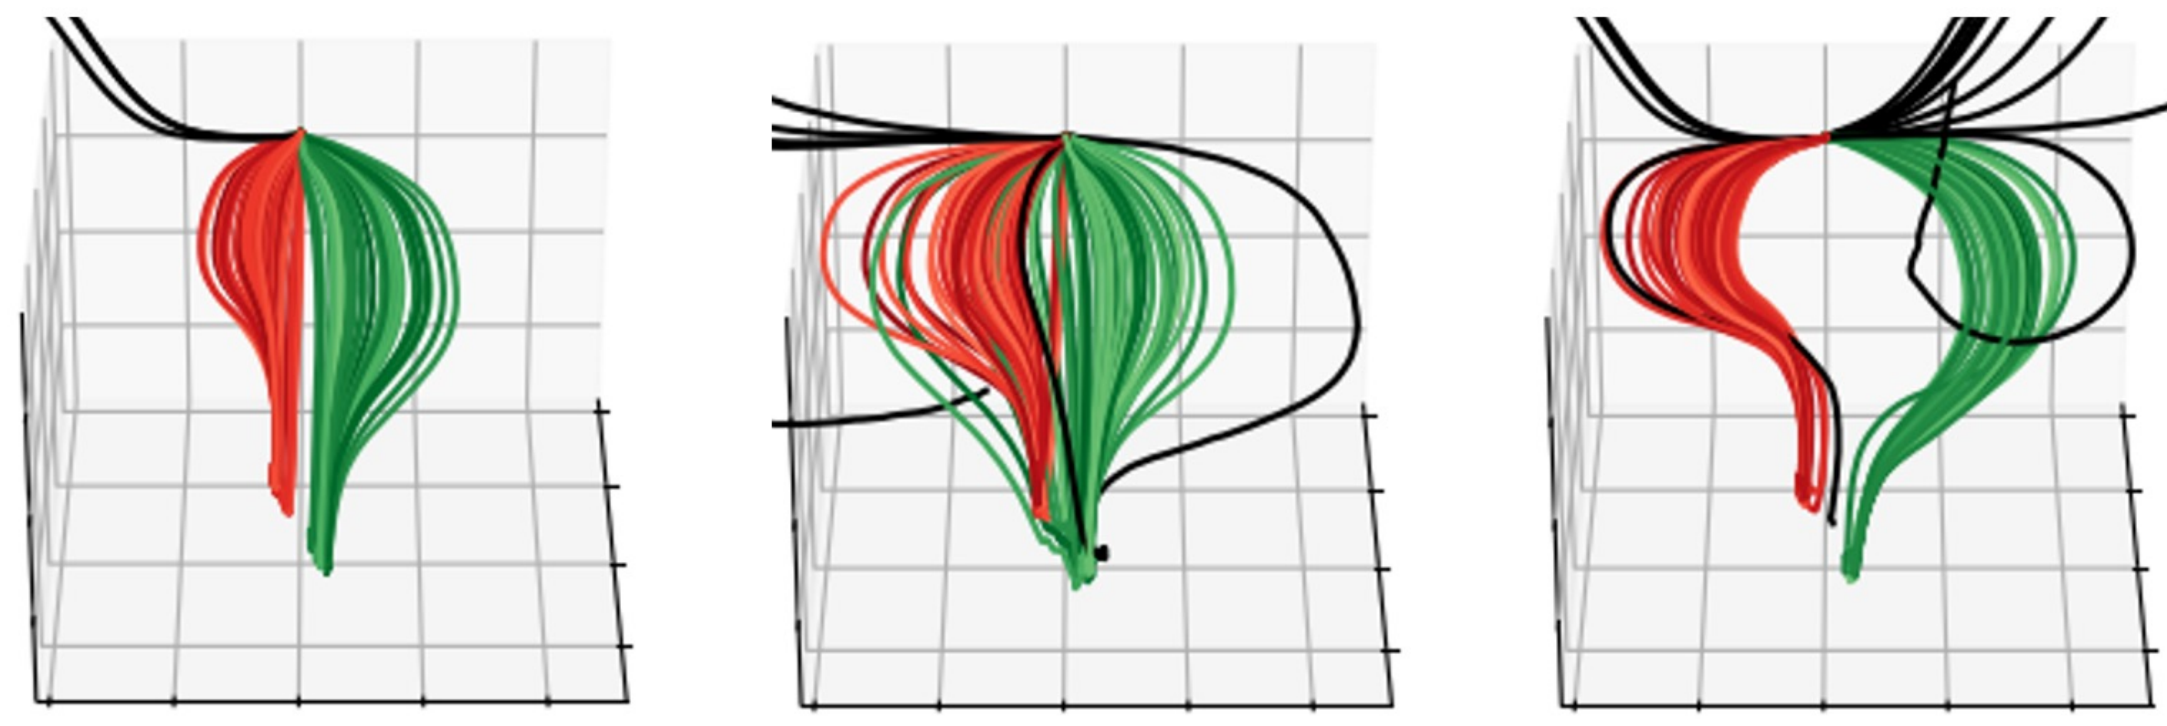


Generative Legible Motion Model



Results

100 Rollouts of Trained Agents




VAE 97% Success G-VAE 93% Success GLMM 86% Success

— Picked Up Red — Picked Up Green — Failure

Evaluating Legibility [3]

$$L(\xi_{S \rightarrow G^*}) = \sum_{s_i \in \xi_{S \rightarrow G^*}} \sum_{\{g \in G: g \neq g^*\}} \frac{\|g - s_i\|_2}{i}$$

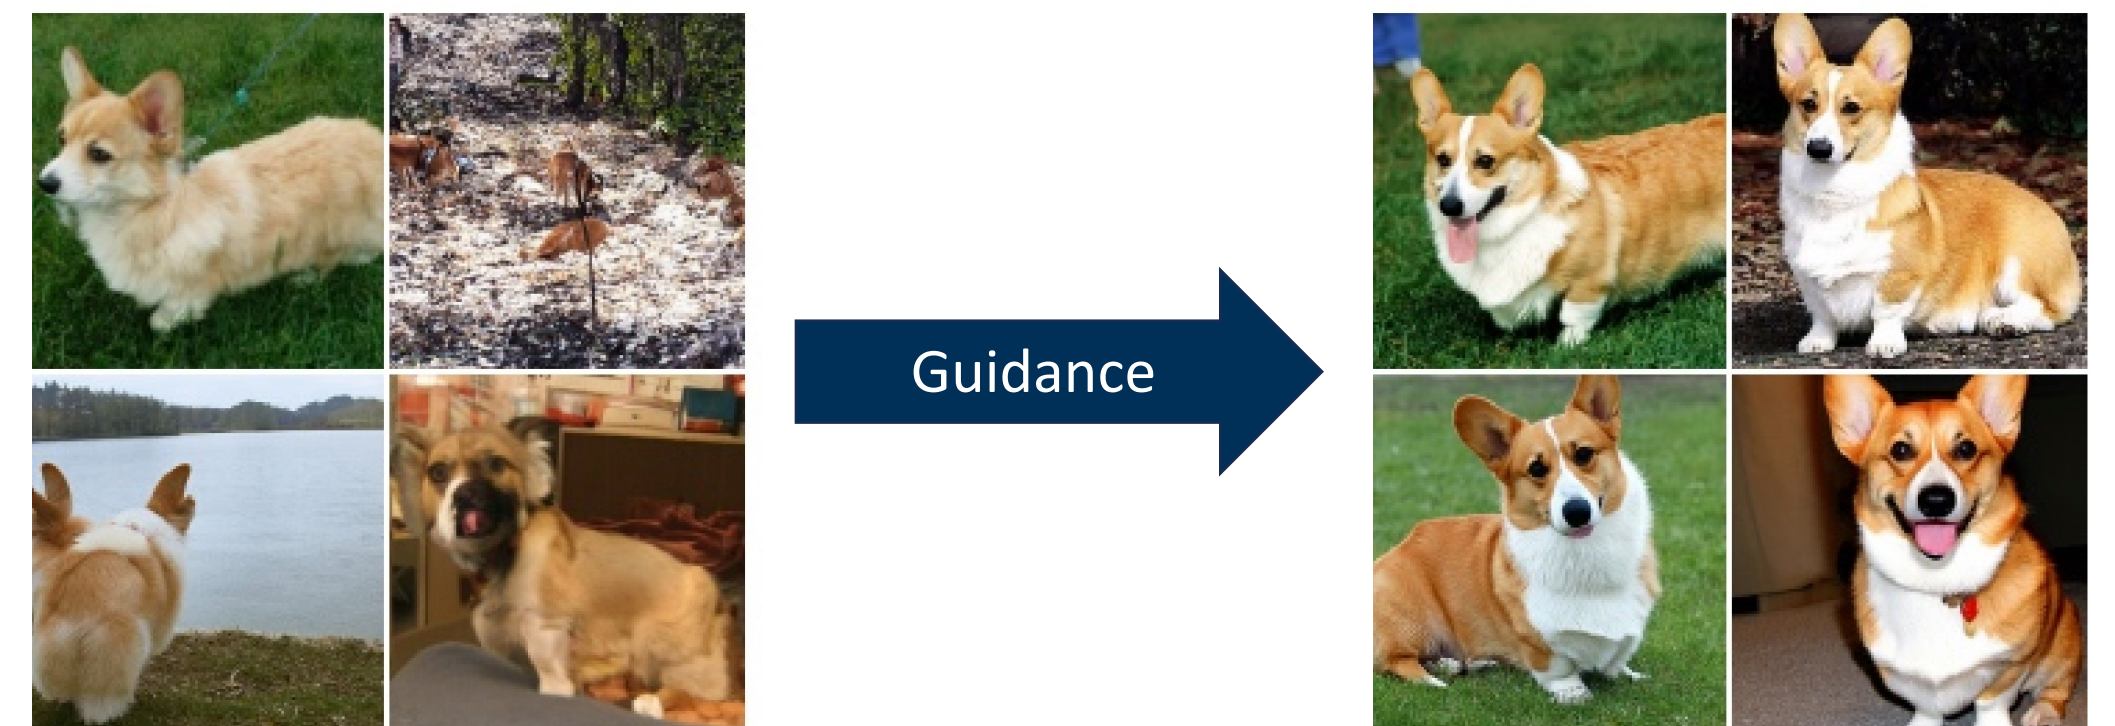
■ VAE ■ G-VAE ■ GLMM



Conclusions

The implicit classifier guidance afforded by GLMM optimizes for legibility, but at the cost of success rate.

Parallels to Image Generation



Unguided Corgi Image Generation

Generation with Classifier-Free Guidance [4]

Image Guidance	Enhances class specific characteristics Improves quality at the cost of diversity
Action Guidance	Selects goal specific states Improves legibility at the cost of success rate

Next Steps

- Implement a legibility classifier conditioned on continuous goal states
- Measure performance of diffusion models with guidance
- Run experiment on more complex tasks
- Directly evaluate how GLMM affects an observer's ability to predict the robot's goal state

References

- [1] Dragan, A. D., Lee, K. C., and Srinivasa, S. S. Legibility and predictability of robot motion. In 2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pp. 301–308. IEEE, 2013.
- [2] Ravichandar, H., Polydoros, A. S., Chernova, S., and Billard, A. Recent advances in robot learning from demonstration. Annual review of control, robotics, and autonomous systems, 3:297–330, 2020.
- [3] Zhao, X., Fan, T., Wang, D., Hu, Z., Han, T., and Pan, J. An actor-critic approach for legible robot motion planner. In 2020 IEEE International Conference on Robotics and Automation (ICRA), pp. 5949–5955. IEEE, 2020.
- [4] Ho, J. and Salimans, T. Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598, 2022.