Closing the Sim-to-Real Gap for Ultra-Low-Cost, Resource-Constrained, Quadruped Robot Platforms

RSS Sim2Real Workshop 2022

Jason Jabbour[1], Sabrina Neuman[2], Mark Mazumder[2], Colby Banbury[2], Shvetank Prakash[2], Brian Plancher[2], Vijay Janapa Reddi[2]

University of Virginia[1]   Harvard University[2]

Motivation

- Imitation learning pipelines have successfully achieved skills on sophisticated robot platforms
- Imitation learning pipelines rely on precise actuators, a menu of sensors, and high maneuverability
- Cannot be directly applied to ultra-low-cost, resource constrained robots

	Unitree A1	Petoi Bittle	Ratio
Cost	$10,000 USD	$299 USD	33x
Weight	12 kg	.29 kg	41x
Dimensions	.5 x .3 x .4 m	.2 x .11 x .11 m	2.5x
Degrees of Freedom (DoF)	12 (Leg: 3)	8 (Leg: 2)	1.5x
Battery Capacity	25.2V 4200mAh	7.4V 1000mAh	3x
Motor Resolution	.022°	1°	45x
IMU	Yes	Yes	NA
Motor Feedback	Yes	No	NA
Foot Pressure Sensor	Yes	No	NA
LiDAR	Yes	No	NA
Computing	ARM Cortex-A72 2.5GHz	Nyboard V1 ATMega328P 20MHz	125x
Optional Additional Computing	NVIDIA TX2 1.3GHz	Raspberry Pi Zero 2W 1GHz	1.3x

Specification comparison between the Unitree A1 and the Petoi Bittle quadruped robots

Challenges: Overview

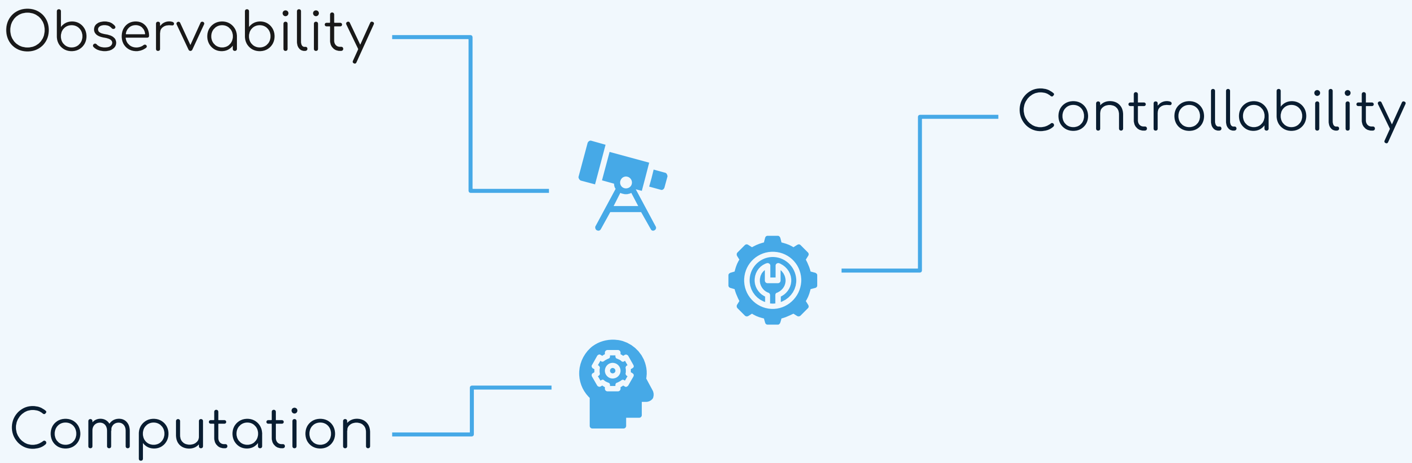

Challenges: Observability

- Bittle's low-cost servo motors do not have encoders to convey their precise position
- Imitation learning pipelines assume perfect knowledge of robot joint positions
- Observational space reduced to accommodate feedback limitations

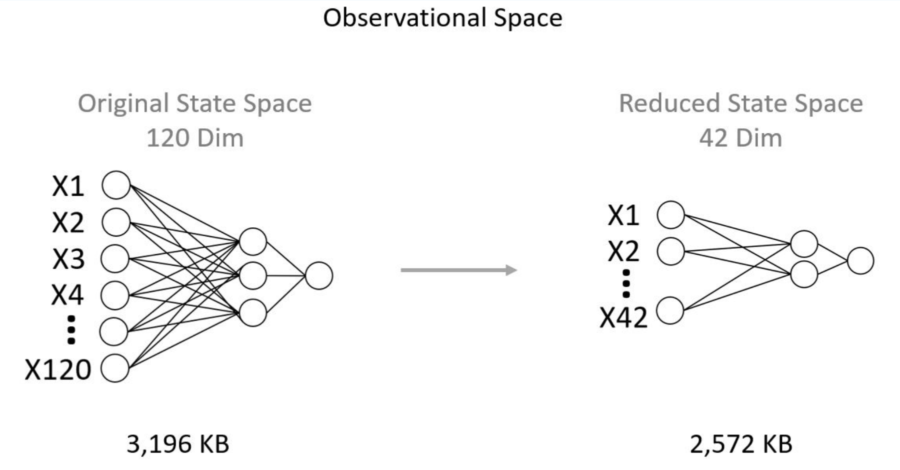

Observational Space Dimension Reduction

Challenges: Computation

- Ultra-low-cost, resource constrained robots have limited onboard computing that prevent running large neural networks

- Applied TinyML techniques, including graph freezing and quantizing model weights to decrease model size

- Decreased model size by a total of **10x** and observed minimal accuracy loss

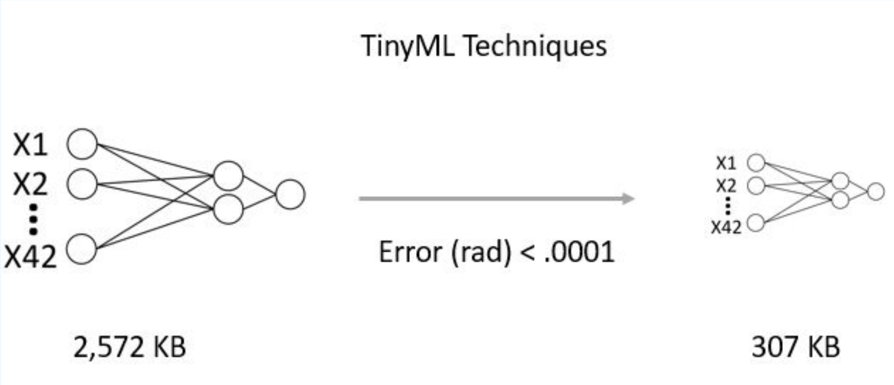

Graph freezing and quantization significantly decreases model size with minimal accuracy loss

Challenges: Controllability

- Imitation learning policies trained for fully-featured robots issue highly precise commands that are unachievable by low-cost actuators
- Incorporated a 1-degree dead-band zone in simulation to reflect real-life hardware limitations
- Policy learns to predict large joint angles that will never be fully achieved
- Predicted actions smoothed to avoid unreasonably large joint angle changes

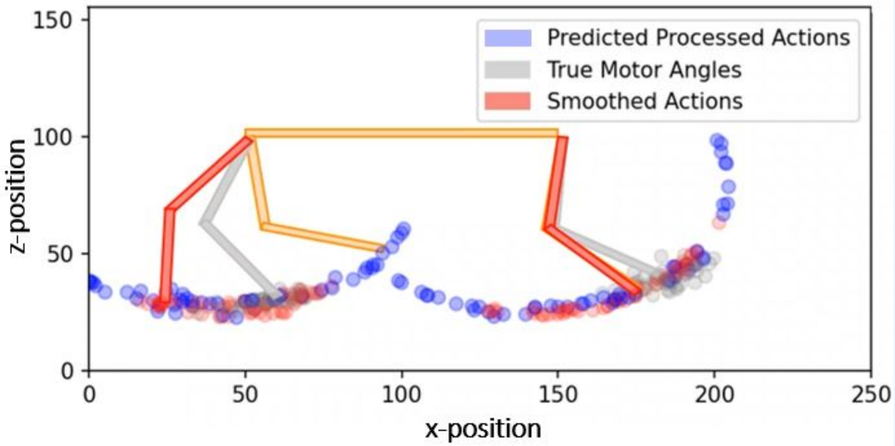

Predicted actions smoothed to accommodate for hardware limitations

Results

Future Work

Hardware Survey and Configuration Layer

Evaluate Policy Sim2Real Transfer Success