SIDDHANT GANGAPURWALA

Date of Birth: 1994 December 07 | **Nationality:** Indian | **Address:** Seattle, USA
Email: siddhant@gangapurwala.com | **Website:** gangapurwala.com

WORK

Sony AI *April 2023 - Present*
Research Scientist in Reinforcement Learning
Location: United States of America

University of Oxford *November 2021 - February 2023*
Postdoctoral Researcher in Machine Learning and Robotics
Research: *Learning Long-Horizon Planning and Control Policies through Inference of World Dynamics*
Lab: Dynamic Robot Systems Group, Oxford Robotics Institute

EDUCATION

University of Oxford *October 2017 - November 2021*
Doctor of Philosophy (DPhil/PhD) in Autonomous Intelligent Machines and Systems
Research: *Learning System-Adaptive Legged Robotic Locomotion Policies*
Supervisors: Dr. Ioannis Havoutis and Prof. Ingmar Posner
Thesis Defence: 2022 January 24 | **Examiners:** Prof. Michiel van de Panne and Prof. Jakob Foerster

University of Mumbai *July 2012 - June 2016*
Bachelor of Engineering (B.E.) in Electronics

INTERESTS

Reinforcement Learning, Robotic Manipulation and Locomotion, Multi-Agent Reinforcement Learning, Simulations and Character Animations

AWARDS

Queen Mary UK Best PhD in Robotics Award *September 2023*

EU Memory of Motion Studentship *2019 - 2021*

EPSRC CDT Autonomous Intelligent Machines and Systems Funding *2017 - 2021*

CONFERENCE PRESENTATIONS AND DEMONSTRATIONS

Conference on Robot Learning (CoRL) *2021*
RLOC: Terrain-Aware Legged Locomotion using Reinforcement Learning and Optimal Control
Demonstration: Live demo of the RLOC framework on the ANYmal C robot

IEEE International Conference on Robotics and Automation (ICRA) *2021*
Real-Time Trajectory Adaptation for Quadrupedal Locomotion using Reinforcement Learning
Presentation: youtu.be/bMtb0raqtaM

NVIDIA GPU Technology Conference (GTC) *2021*
Learning Dynamic and Robust Control Solutions for Robotic Locomotion
Presentation: nvidia.com/en-us/on-demand/session/gtcspring21-s31585

IEEE International Conference on Robotics and Automation (ICRA) *2020*
Guided Constrained Policy Optimization for Dynamics Quadrupedal Robot Locomotion
Presentation: youtu.be/C6n2ZMVxun4

SELECTED PROJECTS

One Policy for Every Quadruped: Platform-Adaptive Robotic Locomotion *2021*
HuboLab, Korea Advanced Institute of Science and Technology

- Trained a dynamics encoding network which maps a history of state transition tuples into a latent representation of the system dynamics.
- Trained a locomotion policy whose behaviour adapts to different quadrupedal platforms using the estimated dynamics.
- Training was performed using domain randomization with procedurally generated quadrupedal models based on a curriculum learning approach.
- Demonstrated multiple trained policies on different platforms in simulation and also on the physical Mini-Cheetah, ANYmal B and A1 quadrupeds.

RLOC: Terrain-Aware Legged Locomotion using Reinforcement Learning and Optimal Control *2020*
*Dynamic Robot Systems Group, Oxford Robotics Institute — **Accepted for publication in T-RO journal***
Manuscript: ieeexplore.ieee.org/document/9779429 | ***Video***: youtu.be/GTI-0gl6Hg0

- Trained a perceptive quadrupedal footstep planning policy to map the proprioceptive and exteroceptive robot state to desired feet positions
- Additionally trained an emergency recovery policy and a domain adaptive tracking policy to adapt to uncertainties in modelled domains.
- Presented a combined reinforcement learning and optimal control based approach to track the neural network parameterized footstep plans using a model-based motion controller.
- Introduced a denoising approach to post-process the local terrain elevation and demonstrated the developed framework on physical systems, ANYmal B and ANYmal C, for locomotion over uneven terrain.

Real-Time Trajectory Adaptation using Reinforcement Learning *2020*
*Dynamic Robot Systems Group, Oxford Robotics Institute — **Published in ICRA 2021***
Manuscript: gangapurwala.com/cltt | ***Video***: youtu.be/Ve4SDl1wI9s

- Presented an approach to perform online replanning of a reference long-horizon motion plan generated using a trajectory optimization solver offline.
- Introduced a procedural terrain generation framework to obtain a policy which generalises to different kinds of environments and trajectories.
- Demonstrated that the trained trajectory-adaptation policy increased the success rate of tracking long-horizon motion plans on a real ANYmal B quadruped even when subject to unexpected perturbations.

Guided Constrained Policy Optimization for Quadrupedal Locomotion *2019*
*Dynamic Robot Systems Group, Oxford Robotics Institute — **Published in RA-L 2020***
Manuscript: gangapurwala.com/gcpo.pdf | ***Video***: youtu.be/iPDmG9knkLs

- Developed a reinforcement learning algorithm for constrained policy optimization such that only policies that strictly obey the necessary safety-critical constraints are sampled for optimization.
- Addressed the issues of sample complexity associated with pure RL strategies by using reference oscillatory motions to warm-start the policy optimization.
- Deployed the trained policy on the real ANYmal B quadruped and demonstrated its performance outdoors on unstructured terrain.
- Additionally demonstrated the robustness of the obtained control policy by emulating a weak knee actuator on a real quadruped, introducing external perturbations, and also changing simulated gravity.

Reinforcement Learning based Solution for Heterogeneous Swarm Optimization *2018*
Robotic Systems Lab, ETH Zürich
Manuscript: gangapurwala.com/hsrl.pdf
Summary: Developed a reinforcement learning training environment to utilise an aerial robot to inspect and map obstacles along the locality of a quadrupedal robot to navigate to a goal while avoiding obstacles.

Generative Adversarial Imitation Learning for Quadrupedal Footstep Planning *2018*
Dynamic Robot Systems Group, Oxford Robotics Institute
Manuscript: gangapurwala.com/gail.pdf | ***Demo***: gangapurwala.com/gtest

TEACHING

Legged Robotic Planning in Unstructured Environments *2020*
Lead Course Instructor - Oxford Robotics Institute, University of Oxford
Led and organized a hands-on course with four instructors for twelve first-year doctoral candidates of the Autonomous Intelligent Machines and Systems CDT.

Robotic Locomotion *2019*
Course Instructor - Oxford Robotics Institute, University of Oxford
Instructor for a hands-on course on robotic locomotion for fourteen first-year doctoral candidates of the Autonomous Intelligent Machines and Systems CDT.

TECHNICAL SKILLS

Programming Languages	C++, Python
Libraries and Frameworks	NVIDIA IsaacGym, Eigen, PyTorch, OpenAI Baselines, ROS
Physics Simulators	RaiSim, NVIDIA IsaacSim, PyBullet, Gazebo

PROJECT ADVISORY

Long-Horizon Motion Planning through Variable Length Sequence Encoding *2022 - Present*
Doctoral Research Project

Learning Robust Control Policies with Minimal Dynamics Randomization *2021 - Present*
Doctoral Research Project

Hierarchical Motion Plan Tracking for a Loco-Manipulation System *2021 - 2023*
Doctoral Research Project

Reactive Manipulation through Learned Environment Occupancy Embedding *2022 - 2023*
Master's Research Project

Learning a Mixture of Navigation and Manipulation Skills for Human Support Robot *2021 - 2022*
Master's Research Project

COLLABORATIONS FOSTERED

Korean Advanced Institute of Science and Technology *2021*
Collaboration on a Research Project between Dynamic Robot Systems group and HuboLab

NVIDIA *2020*
Early Access to NVIDIA IsaacGym framework for Dynamic Robot Systems group

SELECTED PUBLICATIONS

'Learning Low-Frequency Motion Control for Robust and Dynamic Robot Locomotion', **Siddhant Gangapurwala**, Luigi Campanaro and Ioannis Havoutis. IEEE International Conference on Robotics and Automation (ICRA), 2023.

'RLOC: Terrain-Aware Legged Locomotion using Reinforcement Learning and Optimal Control', **Siddhant Gangapurwala**, Mathieu Geisert, Romeo Orsolino, Maurice Fallon and Ioannis Havoutis. IEEE Transactions on Robotics (T-RO), 2022.

'VAE-Loco: Versatile Quadruped Locomotion by Learning a Disentangled Gait Representation', Alexander Mitchell, Wolfgang Merkt, Mathieu Geisert, **Siddhant Gangapurwala**, Martin Engelcke, Oiwi Parker Jones, Ioannis Havoutis and Ingmar Posner. Submitted to IEEE Transactions on Robotics (T-RO), 2022.

'Next Steps: Learning a Disentangled Gait Representation for Versatile Quadruped Locomotion', Alexander Mitchell, Wolfgang Merkt, Mathieu Geisert, **Siddhant Gangapurwala**, Martin Engelcke, Oiwi Parker Jones, Ioannis Havoutis and Ingmar Posner. IEEE International Conference on Robotics and Automation (ICRA), 2022.

'Real-Time Trajectory Adaptation for Quadrupedal Locomotion using Deep Reinforcement Learning', **Siddhant Gangapurwala**, Mathieu Geisert, Romeo Orsolino, Maurice Fallon and Ioannis Havoutis. IEEE International Conference on Robotics and Automation (ICRA), 2021.

'Rapid Stability Margin Estimation for Contact-Rich Locomotion', Romeo Orsolino, **Siddhant Gangapurwala**, Oliwier Melon, Mathieu Geisert, Ioannis Havoutis and Maurice Fallon. IEEE International Conference on Intelligent Robots and Systems (IROS), 2021.

'CPG-ACTOR: Reinforcement Learning for Central Pattern Generators', Luigi Campanaro, **Siddhant Gangapurwala**, Daniele De Martini, Wolfgang Merkt and Ioannis Havoutis. Towards Autonomous Robotic Systems Conference (TAROS), 2021.

'First Steps: Latent-Space Control with Semantic Constraints for Quadruped Locomotion', Alexander Mitchell, Martin Engelcke, Oiwi Parker Jones, David Surovik, **Siddhant Gangapurwala**, Oliwier Melon, Ioannis Havoutis and Ingmar Posner. IEEE International Conference on Intelligent Robots and Systems (IROS), 2020.

'Guided Constrained Policy Optimization for Dynamic Quadrupedal Robot Locomotion', **Siddhant Gangapurwala**, Alexander Mitchell and Ioannis Havoutis. IEEE Robotics and Automation Letters (RA-L), 2020.