AMR 2020/2021: Final Projects

A final project consists of (1) studying papers, notes or documentation on a specific subject (2) performing simulations on a software platform (3) writing a report and giving a presentation (with slides).

Each project must be carried out by a group of **3 students**. To apply, send me an e-mail message with the composition of your group (one message per group, all members in cc:) and a list of **3** projects (identified by numbers) in order of preference. Projects will be assigned on a FIFO basis. 1 or 2-student groups can also apply, but I will merge them into larger groups. The deadline for applying is **December 16**.

Once your group has been assigned a project, we will set up a meeting to discuss the project in detail. You will be assigned a supervisor. When your project is completed, you must send me an e-mail with the report and then I will set up a presentation date. All projects must be completed by **June 30, 2021**.

This is the list of the available projects, followed by a short synopsis (including bibliography) of each project.

FP1. Control strategies for cooperative object transportation by two humanoid robots
FP2. Quadruped gait generation using IS-MPC
FP3. Discrete-time Control Barrier Functions for safe vehicle control
FP4. Optimal trajectory generation in the Duckietown environment
FP5. Safe kinematic control for humanoid robots using CBFs
FP6. Enhancing kinodynamic RRT using CBF-based steering
FP7. Improving footstep planning algorithms by efficient nearest neighbor searching
FP8. Informed asymptotically optimal anytime footstep planning using BIT*
FP9. Strategies for robust gait generation in humanoid robots
FP10. Learning-based NMPC for autonomous race cars
FP11. 3D information-based collision avoidance for safe mobile robot navigation
FP12. Safe robust control by blending RL policies with stochastic MPC
FP13. Risk critic for safe control

---

1. Control strategies for cooperative object transportation by two humanoid robots

It is common for humans to cooperate when transporting heavy and cumbersome objects. This project aims at developing a control framework for replicating this behavior in humanoid robots. In particular, an existing approach for cooperative object transportation must be first implemented and then combined with MPC-based gait generation. The simulation platform will be MATLAB.

- Hawley et al, "Control framework for cooperative object transportation by two humanoid robots", RAS 2019
- Scianca et al, "MPC for humanoid gait generation: Stability and feasibility", T-RO 2020
- Smaldone et al, "Gait generation using intrinsically stable MPC in the presence of persistent disturbances", Humanoids 2019

---

2. Quadruped gait generation using IS-MPC

Controlling the locomotion of legged robots is a complex problem involving nonlinear dynamics and constraints. Simplified models are often adopted to address the problem in real-time with linear optimization-based techniques. The project consists in studying some literature on quadruped locomotion, such as the first paper below, and then developing an algorithm based on the IS-MPC framework to generate a stable gait for a quadruped. The simulation platforms are MATLAB and DART (basic C++ knowledge is preferable).

- Mastalli et al, "Trajectory and foothold optimization using low-dimensional models for rough terrain locomotion", ICRA 2017
- Scianca et al, "MPC for humanoid gait generation: Stability and feasibility", T-RO 2020

3. Discrete-time Control Barrier Functions for safe vehicle control

Control barrier functions (CBFs) have been used to generate safety constraints for robotic systems in the context of optimization-based control. While they have been mainly developed in a continuous-time formulation and then applied in a discretized control scheme, recent works have proposed a discrete-time version of the same concepts. The objective of this project is comparing the discrete-time formulation with the continuous-time one, implementing Nonlinear Model Predictive Control (NMPC) for trajectory tracking on a mobile robot using CBFs for safety constraints (obstacle avoidance/overtaking). The work should cover the scenario proposed in the first paper (i.e. overtaking) but variations are expected (e.g. obstacle avoidance at an intersection). The simulations are to be implemented in Python using the acados solver for NMPC.

- Zeng et al "Safety-critical model predictive control with discrete-time control barrier function," preprint arXiv:2007.11718
- Ames et al "Control barrier functions: Theory and applications," ECC 2019
- https://github.com/acados/acados

4. Optimal trajectory generation in the Duckietown environment

Online trajectory generation poses the challenge of being simple enough to satisfy real-time constraints, yet being able to guarantee performance with respect to the task. An approach consists in reducing the search space by considering the optimization of specific cost functions which yield polynomial trajectories. Thanks to Bellman's optimality principle, this provides temporal consistency (i.e. a replanned trajectory will follow the remaining part of the previously computed solution). The goal of the project is to implement this trajectory generation scheme (coupled with a trajectory tracking controller of choice) to navigate in the Duckietown environment, which consists in a simulator providing customizable city-like scenarios with lanes and intersections. In this setting, the vehicle is controlled using information provided by the monocular camera on top of the vehicle, which is responsible for recognizing the center lane which defines the local Frenet frame. The simulations are to be implemented in the Duckietown simulator (based on OpenAI Gym) using Python, while the software for the camera line detection provided by Duckietown is based on ROS/Python.

- Werling et al. "Optimal trajectory generation for dynamic street scenarios in a Frenet frame," ICRA 2010
- https://docs.duckietown.org/DT19/

5. Safe kinematic control for humanoid robots using CBFs

Differently from standard pseudoinverse-based kinematic controllers, their optimization-based counterpart allows the explicit inclusion of constraints representing hard requirements. The aim of this project is to implement an optimization-based kinematic controller that allows a humanoid robot to realize at best a manipulation task (assigned in terms of a desired trajectory for the position and/or orientation of one hand) while simultaneously satisfying hard requirements such as balance, respect of joint limits, and collision avoidance. The last requirement must be formulated using a Control Barrier Function (CBF), a tool that has been recently introduced for formulating safety-related constraints. The controller should be tested through V-REP simulations in different scenarios using the NAO humanoid.

- Landi et al, "Safety Barrier Functions for Human-Robot Interaction with Industrial Manipulators," ECC 2019

6. Enhancing kinodynamic RRT using CBF-based steering

RRT-based planners, at each iteration, generate a subpath in the configuration space, and extend the tree only if no collision occurs along it. The aim of this project is to implement a RRT-like planner that uses an optimization-based steering method which incorporates an explicit collision avoidance constraint formulated via Control Barrier Functions (CBFs), a tool that has been recently introduced for formulating safety-related constraints. In principle, this allows to guide the generation of subpaths in the free configuration space, thus reducing the number of failed extension attempts and eliminating the need of collision checks. The planner should be tested through V-REP/CoppeliaSim simulations in different scenarios using a nonholonomic system (e.g., a unicycle) and a critical comparison with the basic RRT planner should be carried out.

- Ames at al, "Control barrier functions: Theory and applications," ECC 2019

7. Improving footstep planning algorithms by efficient nearest neighbor searching

One of the bottlenecks in the performance of sampling-based motion planning algorithms is the computational cost of the nearest-neighbor operation. It is, hence, essential to develop efficient techniques for nearest-neighbor searching in order to keep the planning time low. The aim of this project consists in extending the RRT and RRT* footstep planners proposed by our group by implementing a self-balancing k-d tree on non-Euclidean topologies. Simulations on V-REP/CoppeliaSim must be performed to validate the method. This project must be implemented in C++ using ROS and our footstep planning framework.

- Yershova et al, “Improving Motion Planning Algorithms by Efficient Nearest-Neighbor Searching”, T-RO 2007
- Ferrari et al, “An Integrated Motion Planning/Controller for Humanoid Robots on Uneven Ground”, ECC 2019

8. Informed asymptotically optimal anytime footstep planning using BIT*

Popular motion planning techniques include graph-based searches and sampling-based planners. Graph-based searches (e.g. A*) are efficient but performance is dependent on the chosen approximation, which may result in prohibitively long planning time. Sampling-based planners (e.g. RRT*) approximate the problem domain incrementally until a suitable solution is found. However, approximating the domain in every direction may be inefficient. More recent planning methods try to unify and extend these two approaches by using both sampling and heuristics, alternately approximating and searching the problem domain. The aim of this project consists in developing a generic version of BIT* upon our footstep planning framework. Simulations on V-REP/CoppeliaSim must be performed to validate the method. This project must be implemented in C++ using ROS.

- Gammell et al, “Batch Informed Trees (BIT*): Informed Asymptotically Optimal Anytime Search”, RBRS 2020
- Ferrari et al, “An Integrated Motion Planning/Controller for Humanoid Robots on Uneven Ground”, ECC 2019

9. Strategies for robust gait generation in humanoid robots

Humanoid robots can be subject to unpredictable effects, such as collisions with the environment, unevenness of the ground, or simply the effect of unmodeled dynamics. Robust gait generation aims at producing walking gaits that resist to such effects, as long as some bound on the disturbance is available. Villa and Wieber (2017) proposed a robust gait generation scheme based on computing a robust positive invariant set. The goal of the project is to replicate the results of Villa and Wieber (2017), and compare them to IS-MPC with ZMP constraint restriction (Smaldone et al., 2020). The comparison will be done in MATLAB, for which an implementation of IS-MPC will be provided. Dynamic simulations in DART might be considered to extend the project.

- Villa and Wieber, "Model Predictive Control of Biped Walking with Bounded Uncertainties", Humanoids 2017
- Scianca et al., "MPC for Humanoid Gait Generation:Stability and Feasibility", T-RO 2020
- Smaldone et al., "ZMP Constraint Restriction for Robust Gait Generation in Humanoids", ICRA 2020

10. Learning-based NMPC for autonomous race cars

In autonomous racing cars, dynamic effects are especially difficult to model and estimate correctly at high speed. Hewing et al presented an online learning-based strategy for Nonlinear Model Predictive Control (NMPC), in which the unmodeled dynamics of the car is estimated with increasing accuracy at each lap and, in addition, an adaptive robustification of the constraints is applied to prevent their violations. The students should implement the above technique (in MATLAB or Python) and provide a simulation in V-REP/CoppeliaSim.

- Hewing et al, “Cautious NMPC with Gaussian Process Dynamics for Autonomous Miniature Race Cars,” ECC 2017

11. 3D information-based collision avoidance for safe mobile robot navigation

Recently, Deep Learning methods have achieved better-than-human performances in image analysis tasks. Less work has been conducted in the classification and segmentation of objects directly from 3D information such as point clouds or mesh, employing deep neural networks. In this project we aim to implement and test recently published deep learning networks for object segmentation and classification for collision avoidance during a navigation task of a simple mobile robot. The project requires good knowledge of Python.

- https://paperswithcode.com/paper/pointnet-deep-learning-on-point-sets-for-3d
- https://paperswithcode.com/paper/pointconv-deep-convolutional-networks-on-3d
- https://paperswithcode.com/paper/meshcnn-a-network-with-an-edge

12. Safe robust control by blending RL policies with stochastic MPC

Our group recently published a work in which we blended nonlinear Model Predictive Control (MPC) and Reinforcement Learning (RL) policies for safe control. The objective of this project is to extend this previous work to take into account model uncertainties. In particular, a robust optimal control should be implemented based on scenario stochastic MPC, in which each scenario is associated with a set of model parameter sets. After each iteration the MPC scenarios are updated by selecting only the most successful and updating a parameter Gaussian Distribution. Each scenario is associated with a RL policy that allows the solution of the nonlinear MPC. The algorithm will be tested on goal reaching task with obstacles on a unicycle for which only nominal parameters an are known. The project requires good knowledge of Python.

- Turrisi et al, “Enforcing Constraints over Learned Policies via Nonlinear MPC: Application to the Pendubot”, IFAC WC 2020

13. Risk critic for safe control

Although Reinforcement Learning (RL) methods have achieved impressive results in many applications, safety remains a critical bottleneck when deploying these systems for real world problems. This project aims at implementing safe RL policies which are based on a safety critic, a novel concept to tackle safety in the field of Reinforcement Learning. In particular, the considered scenario involves a navigation task for a unicycle platform in the presence of workspace obstacles. The project requires good knowledge of Python.

Srinivasan et al, “Learning to be Safe: Deep RL with a Safety Critic”, preprint arXiv:2010.14603