Konstantinos Vasios

Dipl.-Ing Electrical & Computer Engineer (NTUA GR), specialized in robotics & automation, experienced in the development of control software in the automotive and robotics. Currently a PhD Candidate at the University of Essex trying to bring machine intelligence a step closer to the real-world.

This publication was made possible by the sponsorship and support of Lloyd's Register Foundation (LRF). The work was enabled through, and undertaken at, the National Structural Integrity Research Centre (NSIRC), a postgraduate engineering facility for industry-led research into structural integrity established and managed by TWI through a network of both national and international Universities. Lloyd's Register Foundation helps to protect life and property by supporting engineering-related education, public engagement and the application of research.

Bridging the simulation to reality gap in robotics

Dr. Panos Chatzakos[1], Dr. Mohan Vishwanathan[2], Dr. Jarchi Delaram[2] ,
[1]TWI, [2]University of Essex
2nd Year of PhD

Keywords: sim2real; simulation; robotics; mushroom-harvesting; soft-gripper; soft-robotics; reinforcement-learning;

I. INTRODUCTION

Training, evaluating, and deploying AI agents directly into the real world is most often prohibited by severe time, cost, availability, and safety constraints. Modelling and Simulation are therefore instrumental for the design of AI agents with interesting behavioural characteristics, within the realm of Machine & Reinforcement Learning. We test the efficacy of our approaches against a specific and still practically unsolved manipulation problem, that of robotic mushroom harvesting with a soft robotic gripper. We start with the design of a simulation environment for this specific manipulation task and we proceed with an exploration of sim2real techniques for our real-world setup. We propose a novel control pipeline that encapsulates the simulation, not only as a training platform but also as a live, in-situ, abstraction layer.

II. DESIGN/METHODOLOGY/APPROACH

With the design of the simulation framework for mushroom harvesting immediately two well-known problems arise. Firstly, the complexity and intricacies of the contact dynamics are usually not captured by the commonly utilized rigid multibody physics engines for robotics, rendering a robust sim2real transfer unlikely. Secondly, capturing the behaviour of materials with deformable structures and respective failure modes is a hard problem that is usually framed within analytical Finite Element Analysis (FEM) approaches, which is impractical for our current project objectives.

We choose to design our simulation environment around the PyPullet physics engine, which is commonly used in robotics and reinforcement learning for rigid-multi-body simulations. This system architecture allows us to capture the dominant dynamics of our scene, allows for fast development cycles and lastly it poses an interesting research question, on whether such a degree of system approximation and abstraction suffices for robust sim2real transfer, given the application of explicit sim2real techniques.

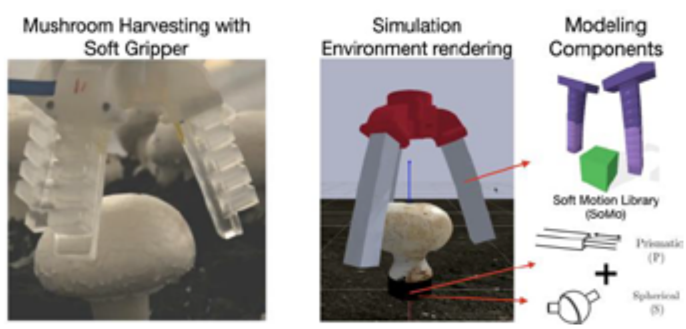

Figure 1 Real soft gripper prototype in the mushroom crop and the equivalent simulation environment.

We capture our scene's dominant dynamics through continuum mechanics formulations and multiple link-joint sequences for the soft robotic gripper, inspired by the Soft Motion Toolkit (SoMo)[1] For the mushroom-root deformation and material failure mode we combine a spherical with a prismatic joint controlled by a simple PD controller for emulating the linear elastic regions(see Figure 1). The proportional gains for all 4 DOFs of the mushroom-root system are computed based on the properties of a liner isotropic material. The failure mode occurs based on the Von-Mises stress criterion.

For fine-tuning the simulation and control parameters we conduct force-position experiments on real-world mushrooms, for determining the material elastic modulus terms and we also conduct step-response experiments on the finger of the real soft-gripper for system identification.

Figure 2 The simulation environment with the Panda Gripper and a randomized scene rendering.

The scene objects textures and colours as well as the scene lighting can arbitrarily change (see Figure 2) for avoiding over-fitting during training, thus increasing generalizing properties and ultimately making a sim2real transfer more likely.

III. FINDINGS/RESULTS

Based on real-world measurements of mushroom stiffness characteristics (see Figure 3) and the soft-gripper step responses we can conclude that our simulation framework is able to capture some dominant & critical dynamics of the mushroom-root system.

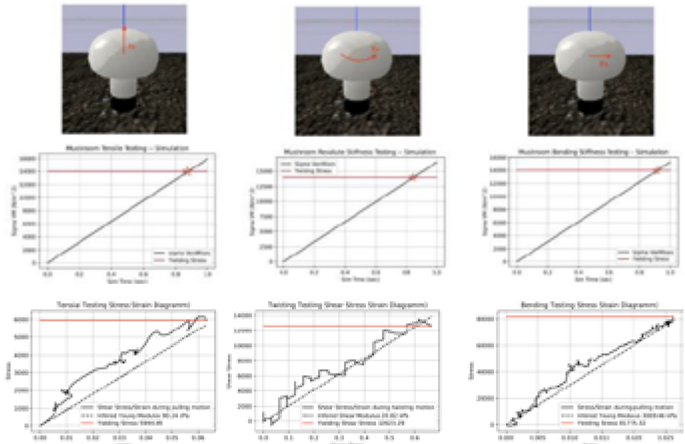

Figure 3 Mushroom root stiffness identification and replication in the simulation environment for the tensile, rotational and bending loads on the mushroom cap.

A very early naive sim2real transfer, without the use of any explicit technique, has been attempted without success, always leading to instabilities and abrupt motions of the real-world robot.

IV. DISCUSSION/CONCLUSIONS

We design a simulation framework for training and testing as a way to circumvent an analytical rule-based planning & control design approach, which would be probably unable to optimally handle each and every corner case for such a complex task (also known as the long-tail distribution problem). Designing a high-fidelity simulation environment for a real-world task, can also potentially be an equally or more complex & prohibitively expensive problem. A viable way to proceed is with a careful selection of the system's critical dominant dynamics and approximation hypotheses together with a set of explicit sim2real techniques.

It is important to note that the design of even a system with simplified-approximate dynamics is not trivial and requires a significant investment on development and some form of system identification processes as well. This is a significant bottleneck for the development and deployment of AI agents for real-world problems.

Another key observation is that the current workflow utilizes the simulation framework only during the training phase. As the simulation framework can be a representation and an abstraction layer of the real-work environment, which is also intuitive for humans, it could be utilized and be an integral part of the perception pipeline.

V. FUTURE PLAN/DIRECTION

Our work now focuses on refining and evaluation the domain randomization as a sim2real technique for this particular problem. For mitigating the problem of the discrepancy between the simulation rendering and the real-world video stream, we are evaluating domain adaptation techniques such as, feature mapping on the CNN embeddings and translation modules for homogenization of the CNN embeddings, regardless of a synthetic or real vision feed.

Based on the key insights we focus our research efforts on automating the generation of an ad-hoc simulation environment on physics-enhanced 3D scene reconstruction methods. The existence of such a simulation layer would allow for plan execution and trajectories generation in agent's "imagination" which can also be examined and evaluated by a human supervisor. These generated trajectories, waypoints and task-specific embeddings can be then translated to real-world robot action.

An intriguing possibility, which we are also examining, is the utilization of LLMs and Generative AI for performing this image-video-to-simulation-environment mapping.

REFERENCES

[1] M. A. Graule, T. P. McCarthy, C. B. Teeple, J. Werfel, and R. J. Wood, 'SoMoGym: A Toolkit for Developing and Evaluating Controllers and Reinforcement Learning Algorithms for Soft Robots', *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 4071–4078, Apr. 2022, doi: 10.1109/LRA.2022.3149580.