Learning to Grasp Objects by Touch

Luke Beddow, Helge Wurdemann, Dimitrios Kanoulas

Department of Computer Science

Grasping objects is easy for humans, but hard for robots.

Packing groceries, for example fulfilling online shopping orders, is predominantly done by hand because humans are faster and more reliable than robots. The reason for this is fruits and vegetables represent a particularly challenging group of objects to grasp due to the following characteristics:

Irregular shapes | Varied surfaces | Easy bruising

We propose a robotic gripper design that adapts to object shapes with flexible fingers and a moveable palm.

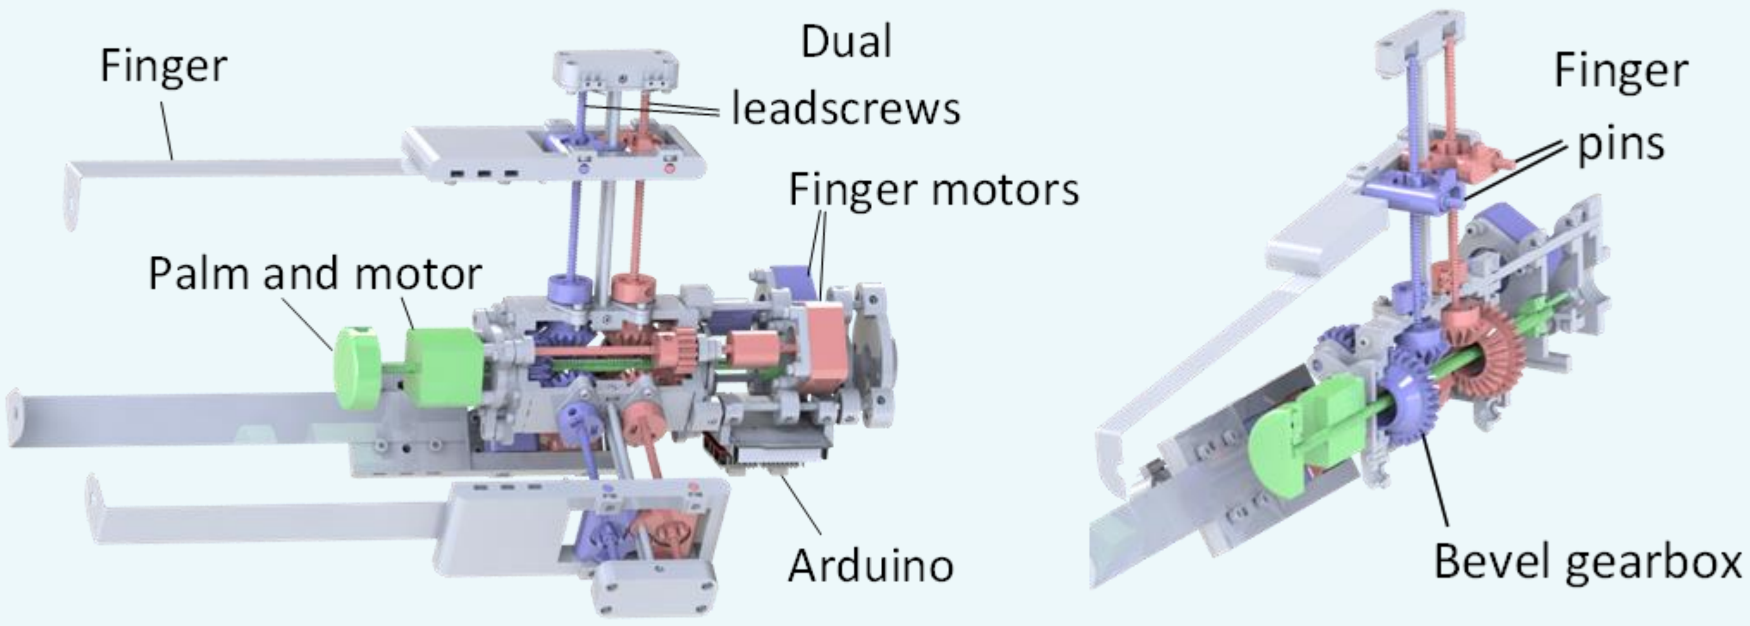

The gripper has three coupled fingers, able to both constrict and tilt. These aim to scoop under objects. Then, the palm is able to move and press the object firmly into the grasp. Grasps were tested with a human, and shown to be stable and to work well on a variety of grocery objects.

Lemon | Bell pepper | Broccoli | Cucumber | Net bag of limes

A machine learning model has been trained to control the gripper by simulating thousands of grasps in a physics engine. It utilises a Deep Q-Network to select actions to complete grasps. Testing is done on an object set of 800 basic object shapes (cubes, cuboids, spheres, cylinders) as most grocery items can be approximated to these. The success rate of the learned policy is 89.6% in simulation; the next step is real world testing.

The bending of the fingers is measured by strain gauges, enabling the robot to feel the grasp. There is also force sensing in the palm and the wrist of the robot (above the gripper) which helps detection objects and the environment. Grasping is completely blind, without any cameras, which means sensor feedback is crucially important.

Several design choices had to be made to enable effective simulated training. The most important was using metal fingers which bend in predictable ways. A mathematical model was derived which split the finger into segments which was then employed in simulation to match the real world behaviour. The below Figure demonstrates the mathematical model performance as the number of segments is increased, for three separate loading cases:

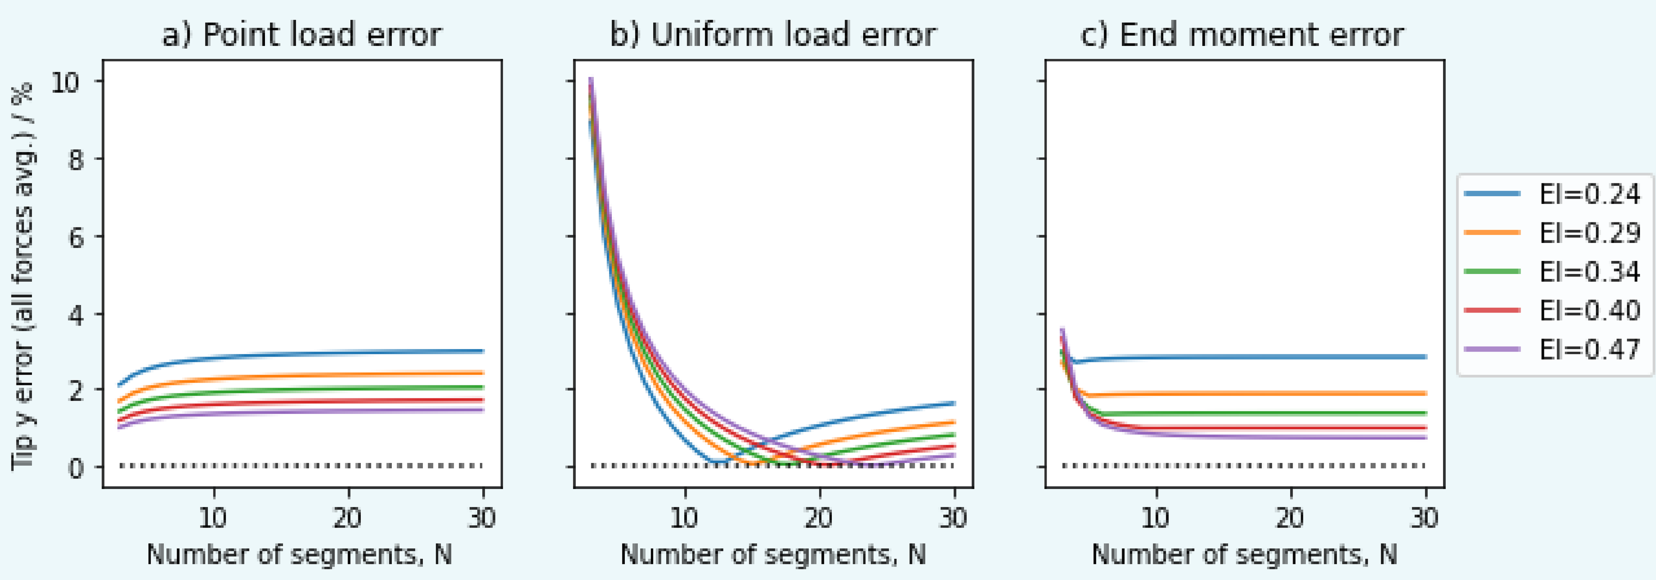

Human operated grasping tests

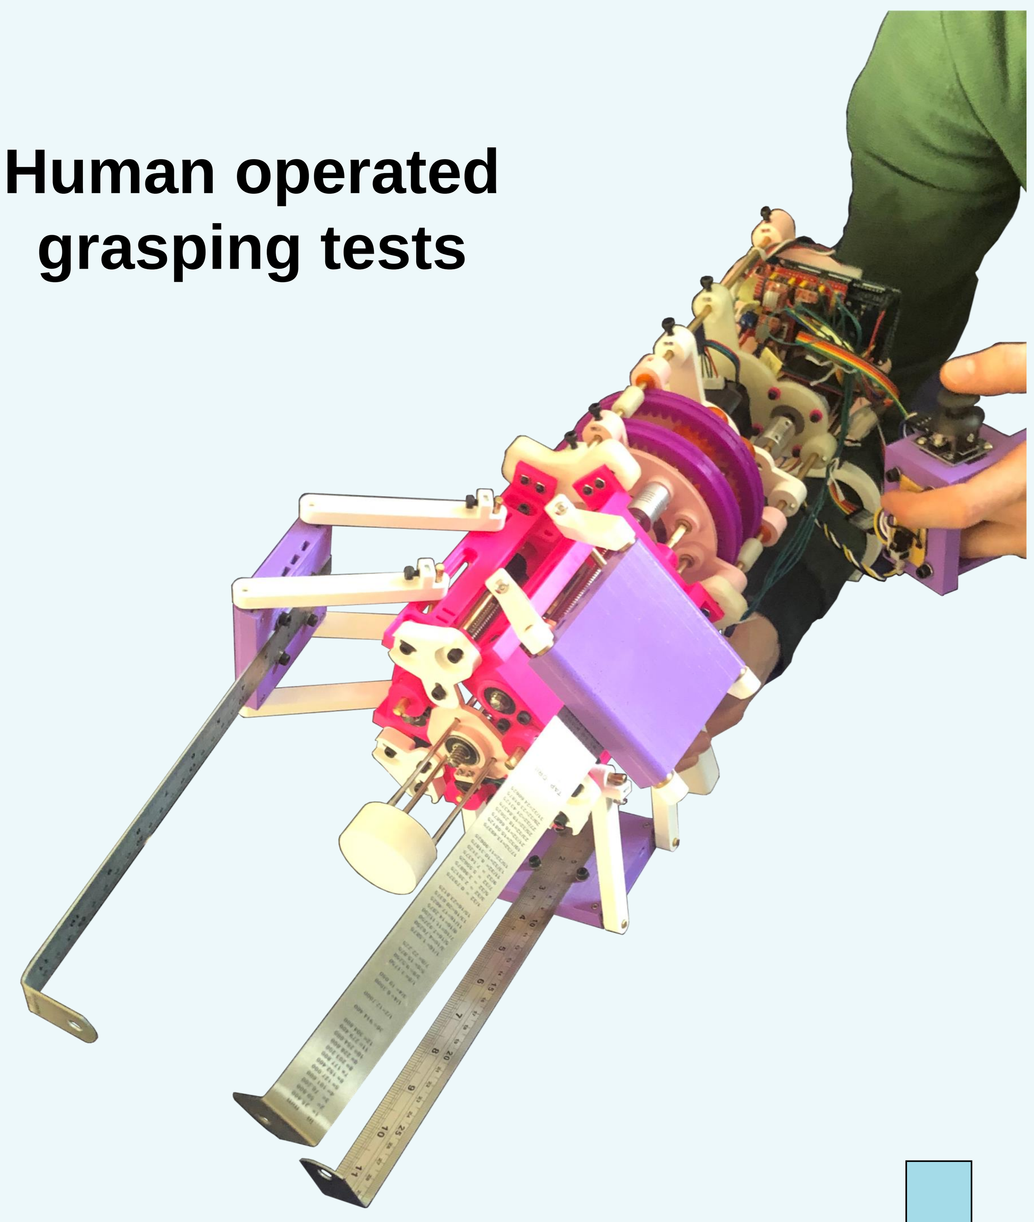

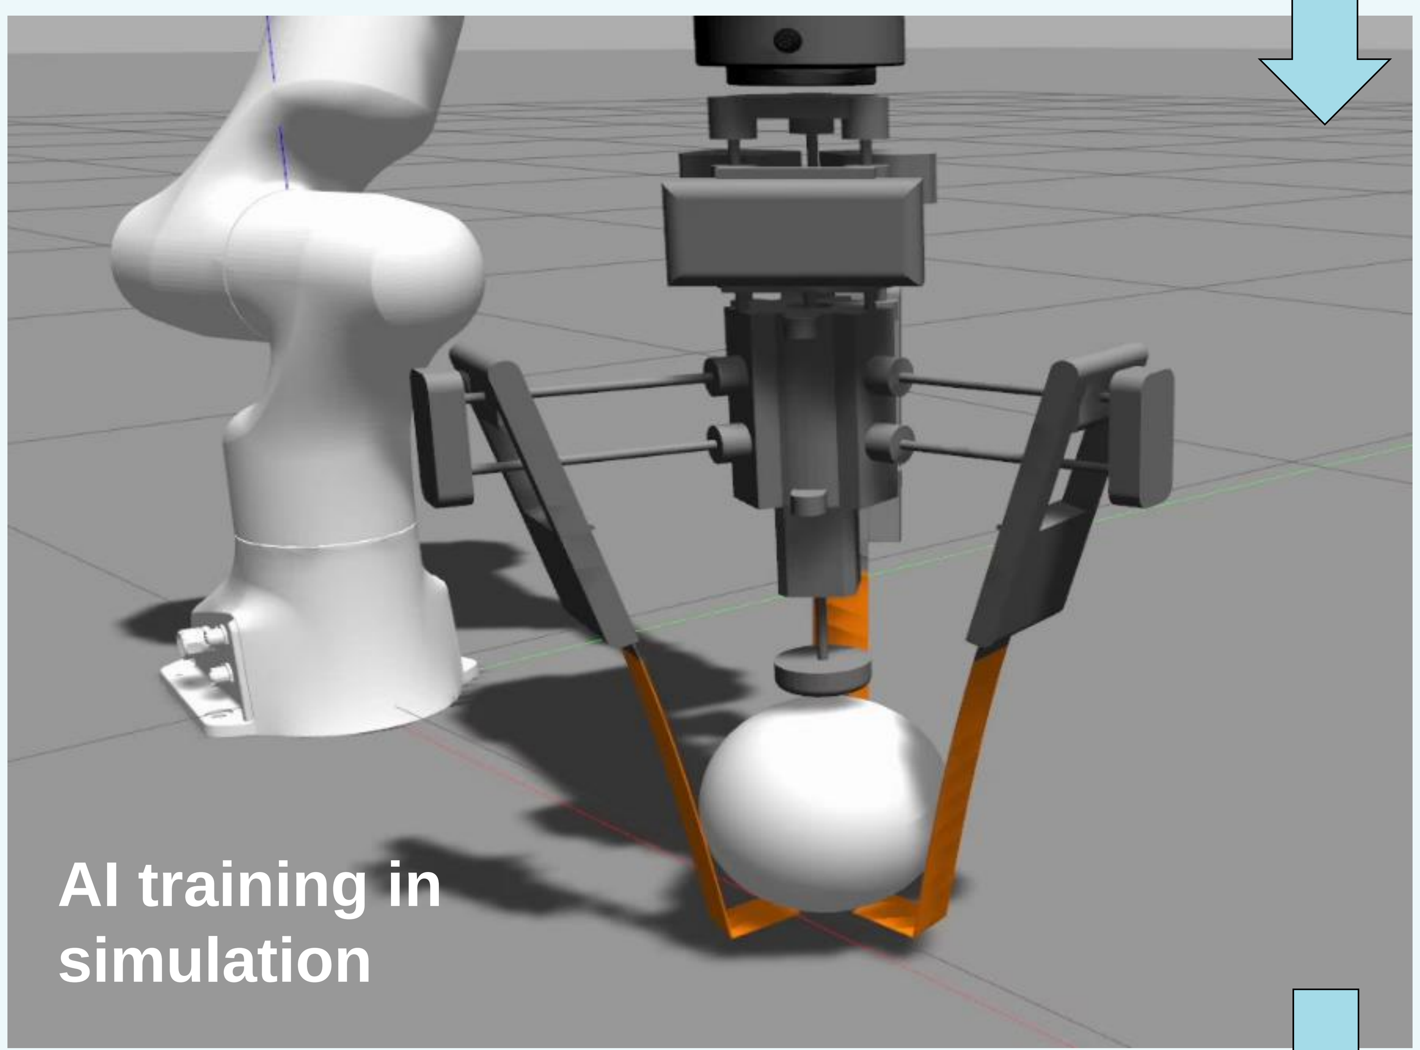

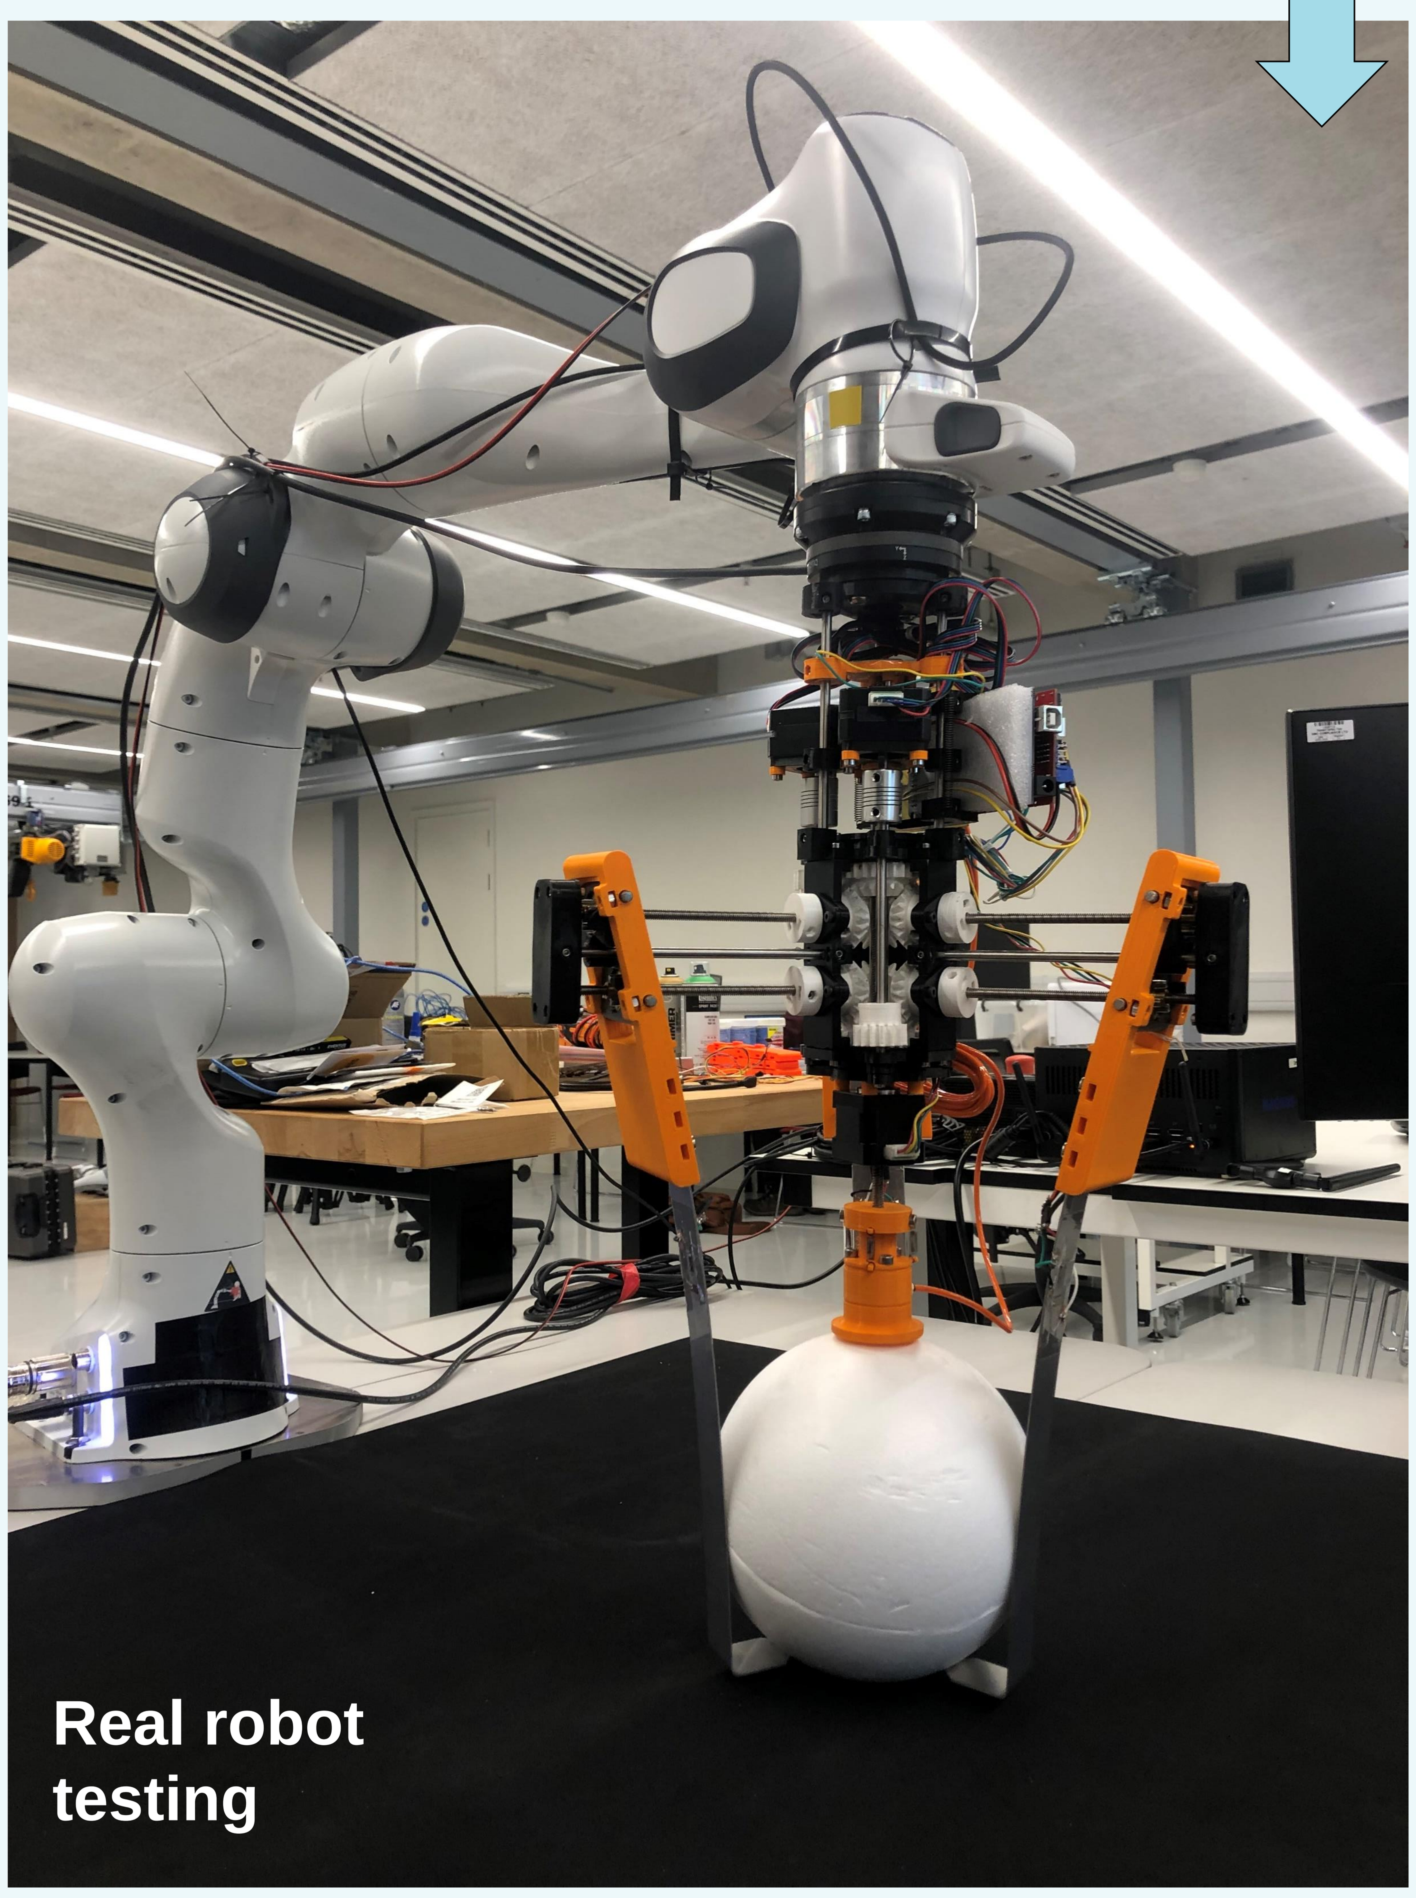

luke.beddow@ucl.ac.uk
h.wurdemann@ucl.ac.uk
d.kanoulas@ucl.ac.uk

Engineering and Physical Sciences Research Council