Learning from Demonstration (LfD) for safe & flexible apple harvesting

Robert van de Ven
Agricultural Biosystems Engineering (WUR)



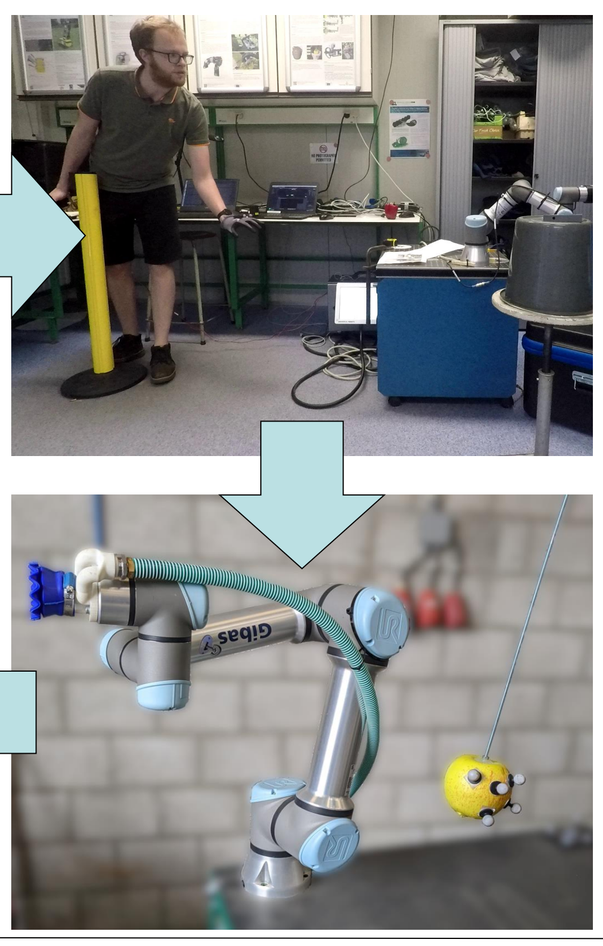
Aims for this course: Learning about learning based control methods for manipulators, connecting with researchers working on control of manipulators for agricultural tasks

Goal: Use LfD to perform agricultural harvesting by incorporating feedback. Instead of pre-programming robot motion plans for each task by robotic experts, LfD offers farmers the opportunity to (re)train a robot for different crops and tasks.

Plans: Four studies, focusing on apple harvesting as an example task:

1. Learning to perform the complete apple harvesting task
2. DualLQR: Approaching oscillating apples efficiently
3. Robust 5D apple pose estimation for correct detachment motion
4. Harvesting apples using LfD in an orchard





Results: Figure 2 shows the structure of the DualLQR controller. Figure 3 shows the performance of the DualLQR controller when tracking an oscillating target in simulation. Here, minimal tracking was observed when far away, combined with precise following when close to the target. Figure 4 shows the result of testing this controller on an apple grasping task. On the right, too tight following resulted in bumping. On the left, too loose following resulted in sliding contact. In the middle, the apple was grasped immediately.

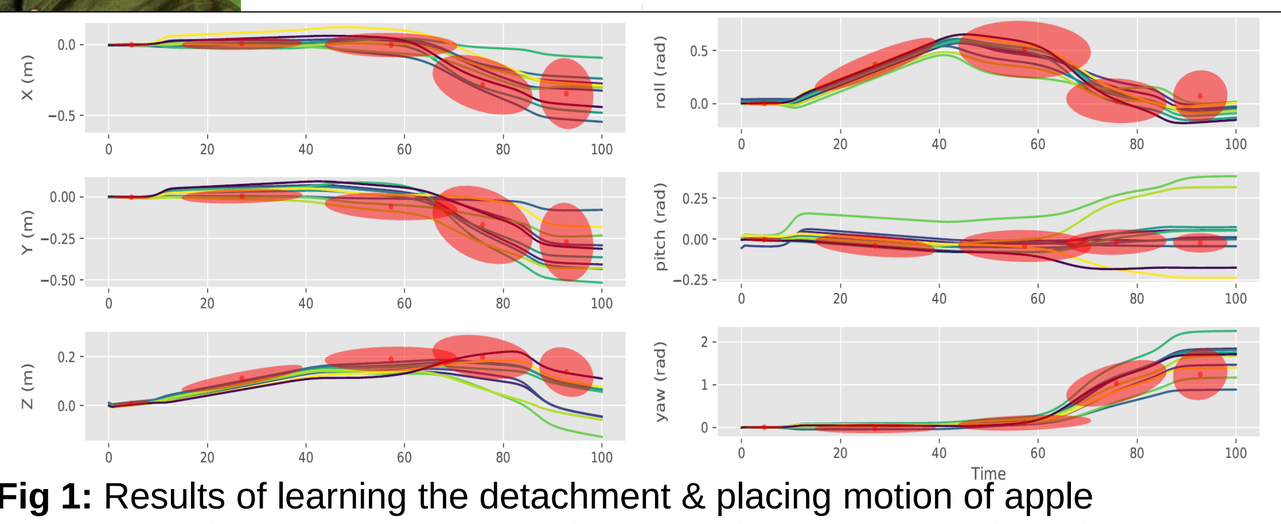
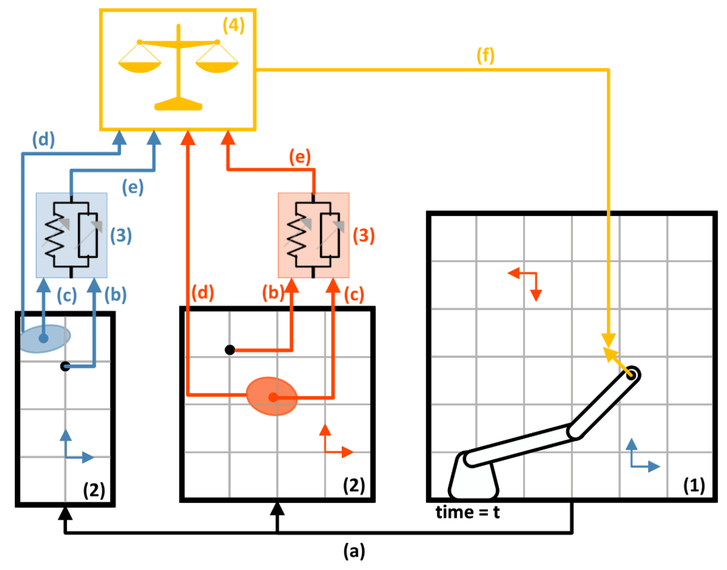


Fig 1: Results of learning the detachment & placing motion of apple harvesting. The orange circles are the learned Gaussians and the lines indicate a path planned by the algorithm.

Results: Figure 1 shows the motion learned using LfD. Here, the specific detachment motion was learned similar to the human demonstrations, with high accuracy in grasp and place pose.

Fig 2: Processes in DualLQR. **(1)** indicates the task in the world frame. **(2)** indicates the task in each reference frame, with the current goal and required precision. **(3)** indicates each LQR. **(4)** indicates the balancing between each LQR. **(b)** indicates the current manipulator pose. **(c)** & **(d)** indicate the current target. **(e)** indicates the control output. **(f)** indicates the combined goal for the manipulator.



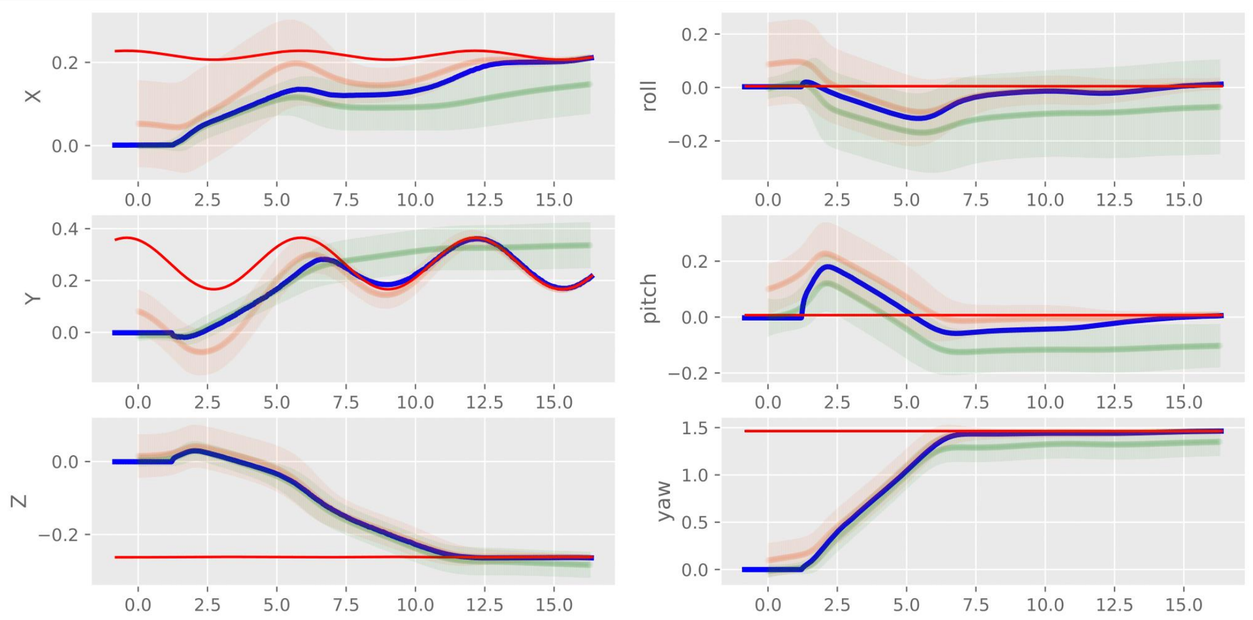
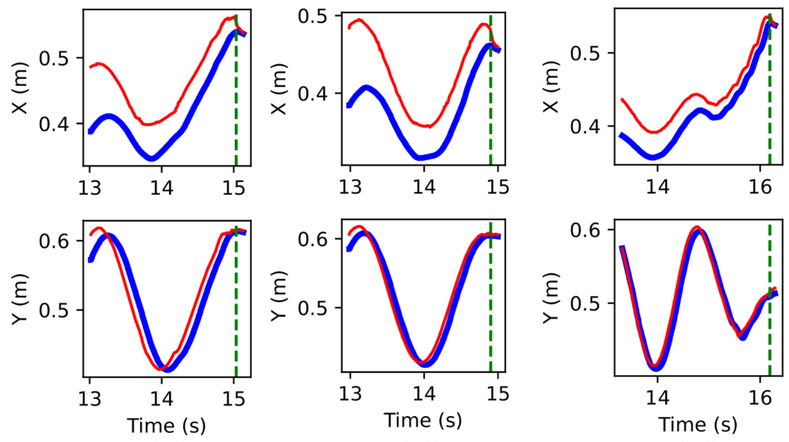


Fig 3: Example of the approach of an oscillating target. The red line indicates the pose of the target over time. The blue line indicates the pose of the manipulator over time. The orange bars indicate the target in the end-frame. The green bars indicate the target in the start-frame.

Fig 4: Trajectory of final approach along X and Y for different controller settings. The red line indicates the pose of the target, the blue line indicates the pose of the robot, and the green vertical line indicates the time of successful grasping.



Control cost: -0.3 -1.2 -3.0

Robert van de Ven
robert.vandeven@wur.nl
ORCID 0009-0005-6483-6231

Image source: https://www.pelckmans.net/nl/blog/fruit-plukken-bewaren-8-tips/details
https://doi.org/10.1016/j.compag.2023.108164

Synergia

WAGENINGEN UNIVERSITY & RESEARCH