Intra-Logistics with Integrated Automatic Deployment:
Safe and Scalable Fleets in Shared Spaces

H2020-ICT-2016-2017
Grant agreement no: 732737

DELIVERABLE 7.3
First demonstration

Due date: month 21 (September 2018)
Deliverable type: DEM
Lead beneficiary: ORU

Dissemination Level: PUBLIC

Main author: Martin Magnusson (ORU)

1 Introduction

This document reports on the outcome of the first physical ILIAD project demonstration on September 27, 2018. The demo coincided with the first stakeholder meeting with invited industry representatives, at the National Centre for Food Manufacturing (NCFM) in Holbeach, UK.

The goal of this milestone demonstration was to show that *"the system can be run in simulation and, to a limited extent, account for real data, execute missions with real vehicles, and/or interact with physical objects (picking and palletizing)"*. As shown in this report, we demonstrated missions executed with a small fleet of two robots in coordination — using live data for calibration, mapping, detection and tracking, and planning — and performed picking of several ILIAD-relevant objects.

A news post was also published shortly after the demonstration on the ILIAD web site at https://iliad-project.eu/second-milestone-demonstration/. This document contains much of the same content.

2 Stakeholder meeting

The milestone demonstration was performed as part of a stakeholder meeting with invited industry represesentatives. Figure 1 shows some impressions from the meeting and demonstrations.

In addition to popular-science presentations of the technical developments in the ILIAD work packages, NCFM presented the project from a food industry perspective. Some key messages from the industry perspective were that inefficiencies in the food and drink industry (which is the largest manufacturing sector in the UK) have significant economic implications, and that employment *increases* in occupations that have been partly automated (because overall demand for their remaining activities has continued to grow). Collaborative research and development, with integrators with a food know-how, will be vital to overcome the current challenges.

During the technical presentations, the audience could already see initial deployment of the ILIAD system (sensor calibration and mapping of the environment). After these presentations, the meeting moved down to the NCFM production facilities, where the audience got a closer view of robot planning and coordination, manipulation, and people perception; as detailed in the next section.

3 Milestone demonstration

In the set of demonstrations, we showcased the current implementation of core components for deployment and operation — from when the first ILIAD robot is unpacked at a new site, to fleet coordination and object detection. Figure 2 shows an overview of the hardware and the environment where the demo was performed.

3.1 Calibration

The very first task when deploying a self-driving warehouse robot should be to calibrate its sensors. The robots in ILIAD do not use any pre-installed infrastructure, but rather use on-board sensors to construct a map of the environment, which is used for localisation and planning. However, in order to do so, each robot must know precisely the position and orientation of its cameras and other sensors. If, for example, a sensor has been shifted from its factory-default mounting during transport, that would affect the reliability and

Figure 1: Images from the first ILIAD stakeholder meeting.

Figure 2: The two specially equipped CiTi-trucks used in the demonstration at the National Centre for Food Manufacturing in Holbeach, UK.

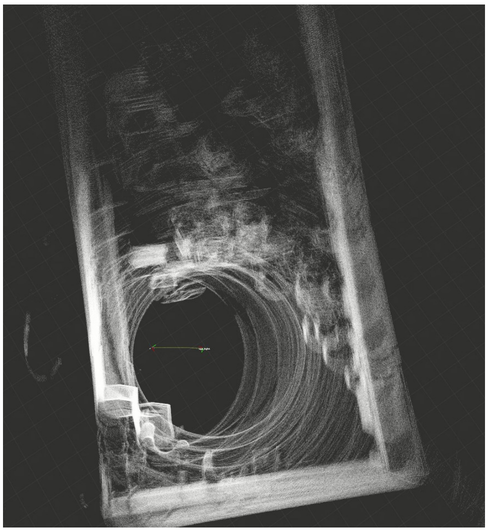

(a) Ten seconds of aggregated lidar data with the initial estimate of the lidar position being at the base of the robot frame, while running the calibration procedure.

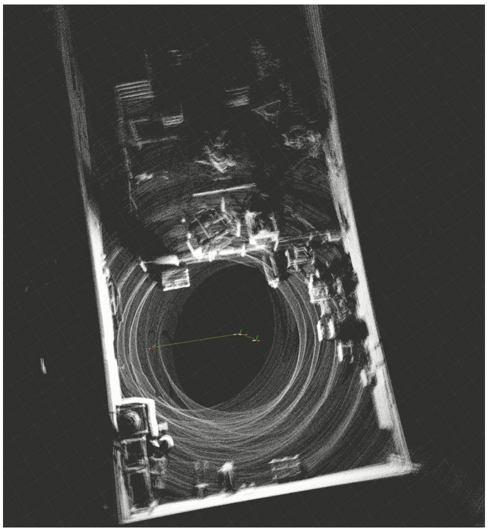

(b) Ten seconds of aggregated lidar data after the calibration is completed.

Figure 3: Automatic calibration routine. These figures show a top-down view of Velodyne 3D lidar scans from the NCFM demo. Figure (a) shows 10 seconds of aggregated data, before calibration. After 16 seconds, the correct pose of the lidar has been found, as shown in (b). Please note that the walls and other features of the room are much more crisp. The remaining blurriness is due to accumulated odometry error over the 10-second time span.

precision of the robot's operation. Therefore, the first task should be to carefully calibrate the sensors on the robot.

As part of ILIAD's Task 1.1, we have implemented a self-calibration routine, which saves considerable time and work during deployment, compared to tedious procedure of standard calibration using custom-made calibration targets. A user only has to walk around with the robot for a few seconds while the calibration software is running, in order to precisely determine the positions of the sensors. The process is visualised in Figure 3. There is also a video available on the ILIAD web site.[1]

The calibration procedure can be qualitatively assessed from the blurriness of aggregated laser scans. When the robot starts moving, without knowing where its sensors are, the room looks very blurry (as in Figure 3a). At 16 seconds into the process, the algorithm finds the correct calibration of the sensor position, after which the image clears up again (Figure 3b).

3.2 Mapping and localisation

Once the calibration is complete, the next step in the deployment phase is to construct a consistent map. Traditionally, localisation of automated trucks (AGVs) is achieved by first installing specific markers in the environment, then manually surveying them, after which they can be used as references to compute the relative position of the truck.

In contrast, the robots in ILIAD are walked through the environment once, during

[1] See http://iliad-project.eu/wp-content/uploads/photo-gallery/MS2/1-calibration-screencast.m4v.

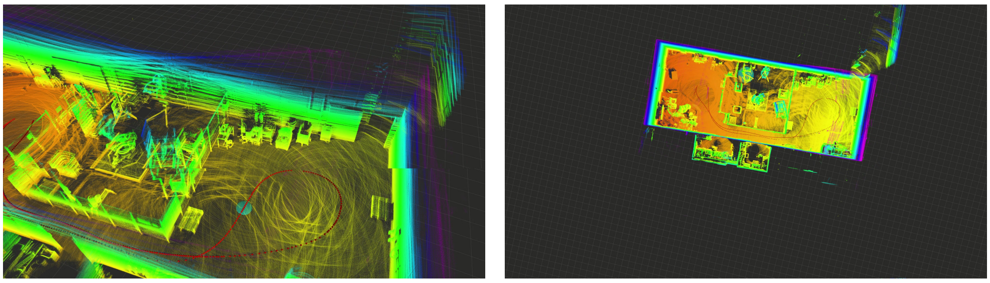

Figure 4: Point cloud map visualisation of the MS2 demo site at NCFM.

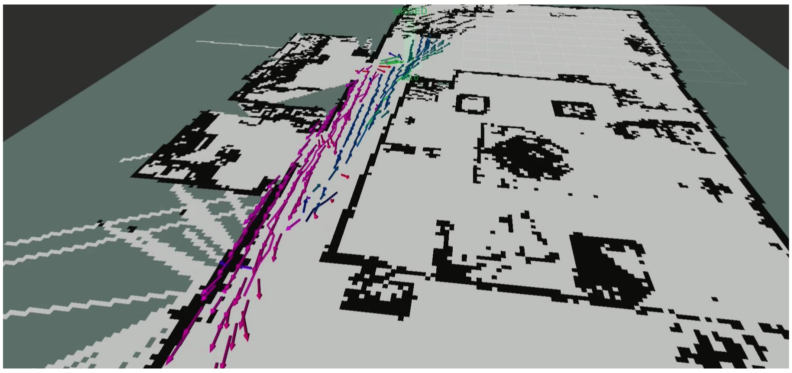

Figure 5: Visualising a CLiFF-map representing learned motion patterns

which they record the shape of the environment, and accurately compensate for any drift that occurs while driving, and at the same time automatically removing moving obstacles from the map, so that only the stationary structures remain in the map. The mapping and localisation procedure was also reported on in Deliverable D1.2.

Figure 4 shows a point-cloud representation of the map created at NCFM. Localisation is performed in a so-called NDT-OM map representation, and 2D motion planning is performed on a 2D grid map extracted from the same data.

We also presented the ongoing work on map representations for learning and prediction of site-specific motion patterns. Figures 5 and 6 show the output of spatio-temporal mapping of the NCFM demo location. Figure 5 shows a representation known as a CLiFF map representing typical motion patterns in and around a corridor, and Figure 6 shows a prediction of the map state at a specific time using the FreMEn representation.

3.3 Task planning

Given an annotation of the map, assigning positions where each product is stored, the fleet is now ready for orders. Assuming that the fleet is connected to a warehouse management system that maintains orders, for each new order, tasks are assigned to the fleet. Figure 7 shows visual output from the pallet planning software provided by partner ACT-OR.[2] Given information about the shapes and weights of each type of package, the system plans how each box should be stacked on the pallet. As of Milestone 2, task allocation from the planned stacking order was not integrated in the system demonstration.

[2]See also video at http://iliad-project.eu/wp-content/uploads/photo-gallery/MS2/3-pallet-planning-and-task-planning.m4v.

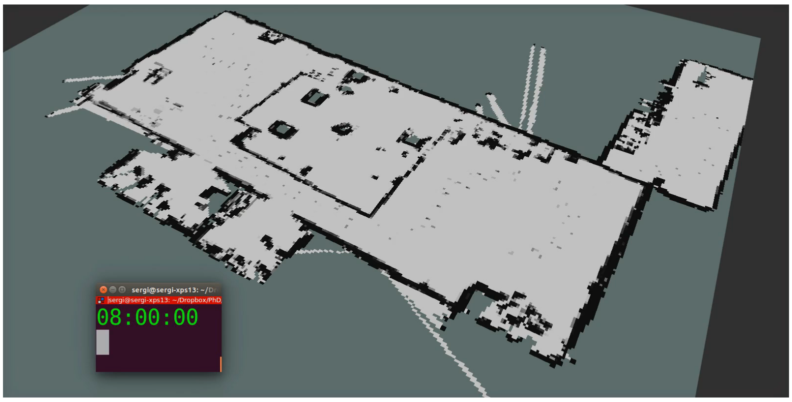

Figure 6: Predicting the state of occupancy for the NCFM environment at a specific time of day.

ORDER DATE AND TIME	ORDER CODE	PRODUCT TYPE	QUANTITY
2017-08-16 13:50	1	Felix Senap short	28
2017-08-16 13:50	1	Felix Senap tall	9
2017-08-16 13:50	1	Felix Tomatpure	9
2017-08-16 13:50	1	Hallonsoppa	4
2017-08-16 13:50	1	Jacky	6
2017-08-16 13:50	1	Onos large	9
2017-08-16 13:50	1	Onos small	30
2017-08-16 13:50	1	Pauluns	12
2017-08-16 13:50	1	Risifrutti	6
2017-08-16 13:50	2	Felix Tomatpure	20
2017-08-16 13:50	2	Hallonsoppa	10
2017-08-16 13:50	2	Jacky	30
2017-08-16 13:50	2	Pauluns	15
2017-08-16 13:50	2	Risifrutti	10
2017-08-16 13:50	3	Hallonsoppa	40
2017-08-16 13:50	3	Jacky	60

(a) Three orders for mixed pallets. The highlighted one (nr 3) was the one used in the simulated test case at Milestone 2.

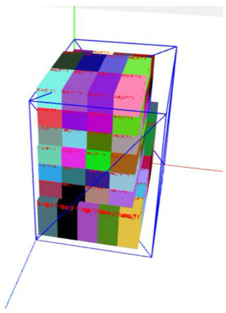

(b) Output from the pallet planning software, with 40 boxes of "Hallonsoppa" raspberry soup and 60 boxes of "Jacky" yoghurt.

Figure 7: ACT-OR's Opt Loading software for task pallet planning, to be used for task assignment.

3.4 Human safety

Detecting and tracking people is a first step towards safe operation (see Section 3.7). However, to make use of that data, it must also be connected to knowledge about humans safety and the control of the vehicle. During the first year of the project, we have built a database of human injury biomechanics. Over 50 years of data from experiments on human injury biomechanics, specifically for the head and chest, have been processed by TUM and distilled into a database that is being used within ILIAD to develop a "safe motion unit" (SMU) for fork-lift trucks and robot arms.

At the MS2 demonstration, a first version of the SMU was demonstrated to control the speed of a robot arm performing a pick-and-place task. This demonstration was performed via video link from Munich. A Franka robot arm was adjusting its joint speeds w. r. t. a "induced human". Full integration of the ILIAD people detection with the SMU is planned for MS3.

3.5 Coordination

Now that each robot knows in which order to put objects on the pallet, the system plans how the fleet of robots should move in order to fulfill the order. Given the map created during the deployment phase (and updated during long-term operation), the robots plan and coordinate their paths on the fly, without the need to manually design paths or traffic rules.

This on-line motion planning and coordination further cuts deployment costs of the system by foregoing the need to carefully craft paths and ad-hoc precedence rules for the robots in the fleet. It also adds to the flexibility, as it makes it possible for robots to replan – in case of unexpected obstacles, for example.

A particular strength of the coordination framework developed as part of ILIAD is that it is very flexible with respect to a heterogeneous fleet of robots. We are able to cater for robots with quite different kinematics and motion capabilities. This makes adding new robots to the fleet quite trivial.

Figure 8 shows an example with a fleet of two robots.[3]

3.6 Manipulation

In this milestone demonstration, we also showed picking of multiple types of objects with a dual-arm manipulator. The arms are not currently integrated with the truck platforms. Instead, the picking was demonstrated via video link from the University of Pisa (UNIPI). Once a robot truck with arms reaches a picking location, this is how it will pick objects and place it on the pallet of the current order. UNIPI have developed a dual-arm manipulator, where one arm has a five-fingered hand and the other has a small tray with a conveyor belt (see Figure 9). Videos from the MS2 object picking demonstration are available online.[4]

There are several challenges for manipulation in warehouses, not least within the food industry which is ILIAD's main use case. While there are several systems commercially available that can pick also unsorted objects rather efficiently, they typically work with suction grippers that pick up each object from above. For many factory/assembly tasks, this is sufficient, but for warehouse objects, the situation is more challenging. One challenge is the variability of objects. Products in warehouses come in many shapes and sizes, and it is difficult to construct a single gripper that can handle all types of objects.

[3] See also http://iliad-project.eu/wp-content/uploads/photo-gallery/MS2/4-fleet-planning-and-coordination.mp4.

[4] See https://iliad-project.eu/wp-content/uploads/photo-gallery/MS2/5-picking-hallonsoppa_row.mp4 and https://iliad-project.eu/wp-content/uploads/photo-gallery/MS2/5-picking-senap_row.mp4.

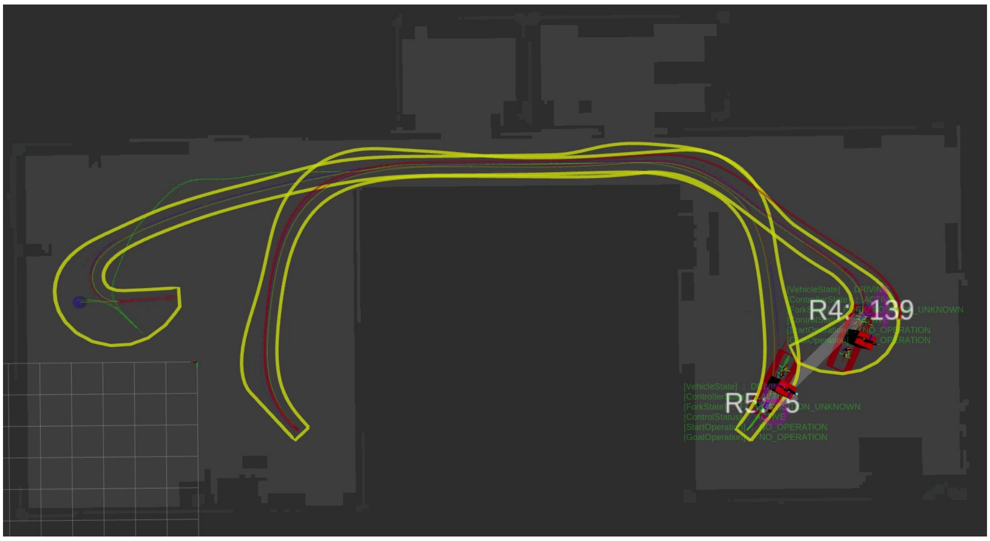

(a) Visualisation of the fleet coordinator. The two trucks have both planned paths to the left side of the hall, but due to the narrow corridor, the trucks must coordinate who goes first. The yellow outlines show the envelopes that the robots will sweep while traversing their paths. The thick grey arrow between the robots indicates that there is a critical point, computed by the coordinator, at which truck R4 will have precedence over truck R5.

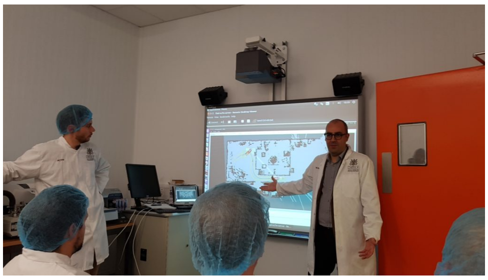

(b) Demonstrating the coordinator demo to the participants of the stakeholder meeting, while the robots are running in the adjacent room.

(c) Overhead view of the demonstration.

Figure 8: Demonstrating the ILIAD fleet coordinator.

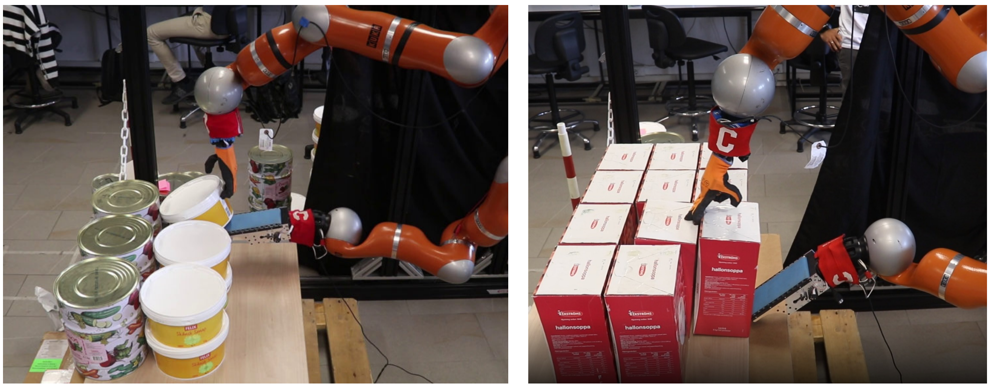

Figure 9: Demonstrating picking of heterogeneous objects: mustard buckets and hallon-soppa (raspberry soup mix).

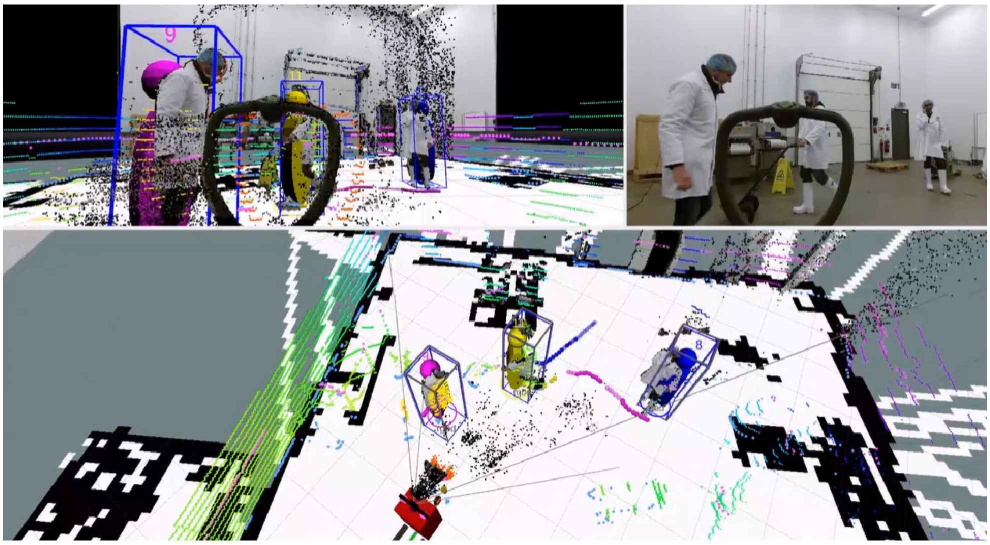

Figure 10: People detection and tracking with RGB-D data at the milestone demonstration. Top right: camera image feed. Top left: overlayed with detections, and point cloud from the Kinect One sensor. Bottom: top view, also showing tracking information.

The ILIAD dual-arm manipulator is much more flexible and human-like, and therefore handles a much larger range of products. Another challenge is that products are heavy, and that the packing is often not designed to be lifted from the top. In food warehouses, many objects are in open-top boxes, which means that they cannot be picked from above. Also the objects shown in Figure 9 could not be picked from above. For the mustard buckets, the lid might come off, and the cardboard boxes shown are heavy enough so that the package may break if it is lifted from above. Challenges such as these, together with the fact that products are stacked tightly, means that a more human-like picking solution is required in order to integrate well with existing warehouse operations.

UNIPI also demonstrated a preliminary version of their system for compliant robotic cutting of plastic stretch wrap. This is a necessary first step before objects can be picked, since pallets are wrapped in stretch wrap to keep objects from falling when they enter the warehouse.

3.7 People detection and tracking

One of the key aspects of ILIAD (in addition to facilitating automatic deployment and adaptation to changing environments) is safe operation among people. One of the cornerstones of this capability is to reliably detect persons in the surroundings of the robot. We demonstrated people detection and tracking software at the MS2 demonstration. Figure 10 shows sample output from the RGB-D detection and tracking modules. This people tracking system will be further robustified using data from additional sensor modalities (lidar and a dedicated safety-vest camera) and demonstrated in the remaining milestone demonstrations.

Finally, we demonstrated the present version of our object detection module, spefically its performance when it comes to detecting pallets (see Figure 11).

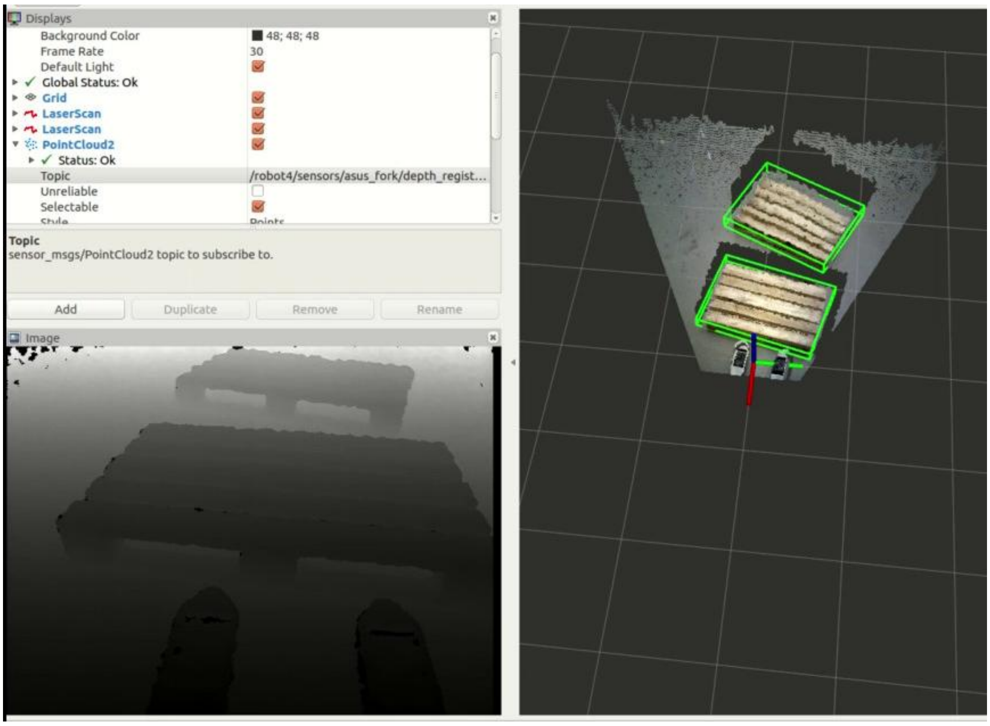

Figure 11: Pallet detection.

4 Summary

In conclusion, this milestone demonstration showed the integration of functional prototypes of some of the most central capabilities of the ILIAD system. In October 2019, at the Milestone 3 demo, we will demonstrate the integration also of the more long-term aspects of ILIAD, this time at a real-world warehouse of food manufacturer Orkla Foods.