

# Towards the deployment of an autonomous last-mile delivery robot in urban areas

A. Santamaria-Navarro<sup>1</sup>, S. Hernández<sup>1</sup>, F. Herrero<sup>1</sup>, A. López<sup>1</sup>, I. del Pino<sup>1,6</sup>, N. Rodríguez<sup>5</sup>, C. Fernández<sup>4</sup>, A. Baldó<sup>2</sup>, C. Lemardelé<sup>2,3</sup>, A. Garrell<sup>1</sup>, J. Vallvé<sup>1</sup>, H. Taher<sup>1</sup>, A.M. Puig-Pey<sup>1</sup>, L. Pagès<sup>2</sup>, A. Sanfeliu<sup>1</sup>

**Abstract**—Nowadays, the skyrocketing last-mile freight transportation in urban areas is leading to very negative effects (e.g., pollution, noise or traffic congestion), which could be minimized by using autonomous electric vehicles. In this sense, this paper presents the first prototype of Ona, an autonomous last-mile delivery robot that, in contrast to existing platforms, has a medium-sized storage capacity with the capability of navigating in both street and pedestrian areas. Here, we describe the platform and position it with respect to other existing prototypes, providing its main Software modules and the first validation experiments, carried out in the Barcelona Robot Lab (Universitat Politècnica de Catalunya); Esplugues de Llobregat (next to Barcelona); and Debrecen (Hungary), which are representative urban scenarios. In such validations, we focus our analysis on the key localization module, whose errors could cascade down the rest of the navigation pipeline (e.g., planning or control). Aside from robotic technical details, we also include the results of the technology acceptance by the public present in the Esplugues de Llobregat test, collected in situ through a survey.

## I. INTRODUCTION

The logistics chain can be divided in two main transportation phases: a) the long-range transportation, where the package travels from one warehouse to another; and b) the short-range operation (referred as last-mile hereafter), where the package is transported from/to the warehouse to/from the customer. While in the long-range phases the focus is on developing efficient transportation systems or global coordination techniques (among others), the short-range operations may require precise maneuvering, interaction with people or sorting semi-structured terrains.

Regarding last-mile operations, there are several factors stressing the city logistics. For instance, urban population growth and rapid urbanization have generated an increase in freight transportation. Moreover, the appearance and rapid

acceptance of e-commerce imply a high amount of goods to be delivered in metropolitan areas, which is expected to increase dramatically in the next few years with more frequent and fragmented deliveries [1], [2]. Today's logistics operations in city centers lead to very negative effects: increase in traffic congestion; air and noise pollution; or safety problems for pedestrians, bikers and deliverers. Researchers and institutions are suggesting a more integrated logistics system, where shippers, carriers, and movements are coordinated, and the freight of different customers and carriers is consolidated into the same *green* vehicles, while avoiding some business models depending on unfair and low-cost labor (e.g., e-riders) [3]. In this sense, this paper presents our first version of such *green* vehicle, an autonomous last-mile delivery device named Ona (shown in Figure 1).

Autonomous city navigation (namely AutoNav) is a hot topic ([4], [5]). In particular, some examples of autonomous last-mile vehicles may include: The Kiwibot<sup>1</sup>, a four-wheeled device designed to deliver very small items; similarly, the Starship<sup>2</sup> platform, meant to drive on sidewalks; the Scout<sup>3</sup>, a small six-wheeled platform developed by Amazon; Serve<sup>4</sup>, by Postmates, which is also designed to carry small items but with a higher ground clearance; Camello<sup>5</sup>, a platform developed by OTSAW and also meant for sidewalk navigation; Roxo<sup>6</sup>, the FedEx SameDay Bot; Clevon 1<sup>7</sup>, with an adaptable, lightweight multi-purpose platform; and Nuro<sup>8</sup>, a high-capacity autonomous robot meant to navigate on the street or in wide pedestrian areas.

All these robot have similar modules related to perception (to gather information of the surroundings); localization and mapping (to estimate the state of the vehicles and represent the environment); planning and decision-making (to generate safe trajectories); control (to drive the platform); and human-robot interaction (to interact with the users or bystanders). However, from all the mentioned (under development) platforms, shown in Figure 2 ordered by scale, Ona's design is closer to Clevon 1 or Nuro. Clevon 1 has been designed as adaptable platform for different business cases related to goods' delivery and, similar to Ona and in contrast to Nuro, favoring the placement of the autonomy payload (se.g., sensors) on the lower level

<sup>1</sup>A. Santamaria-Navarro, S. Hernández, F. Herrero, A. López, I. del Pino, A. Garrell, J. Vallvé, H. Taher, A.M. Puig-Pey and A. Sanfeliu are with the Institut de Robòtica i Informàtica Industrial, CSIC-UPC, C. Llorens i Artigas, 4-6, 08028, Barcelona, Spain. {asantamaria; shernand; fherrero; alopez; agarrell; jvallve; idelpino; htaher; apuigpey; sanfeliu}@iri.upc.edu. <sup>2</sup>A. Baldó, C. Lemardelé and L. Pagès are with CARNET, C. Jordi Girona 1-3, 08034, Barcelona, Spain. {albert.baldo; clement.lemardele; laia.pages}@carnetbarcelona.com. <sup>3</sup>C. Lemardelé is also with the Civil Engineering School of Barcelona, Department of Civil and Environmental Engineering, UPC-Barcelona TECH, C. Jordi Girona 1-3, Building B1, North Campus, 08034 Barcelona, Spain. <sup>4</sup>C. Fernández is with Urbiotica, C/ Joaquim Molins 5, 4t 4a, 08028, Barcelona, Spain. carlos.fernandez@urbiotica.com. <sup>5</sup>N. Rodríguez is with the Instituto de Astrofísica de Canarias, IAC, C. Vía Láctea, 1, 38205, la Laguna, Santa Cruz de Tenerife, Spain. nicolas.linares@iac.es. <sup>6</sup>I. del Pino is also with the department of physics, systems engineering and signal theory of the University of Alicante, Campus de San Vicente del Raspeig, s/n, 03690 San Vicente del Raspeig (Alicante), Spain.

<sup>1</sup><https://kiwibot.com>

<sup>2</sup><https://starship.xyz>

<sup>3</sup>[https://bit.ly/scout\\_amazon](https://bit.ly/scout_amazon)

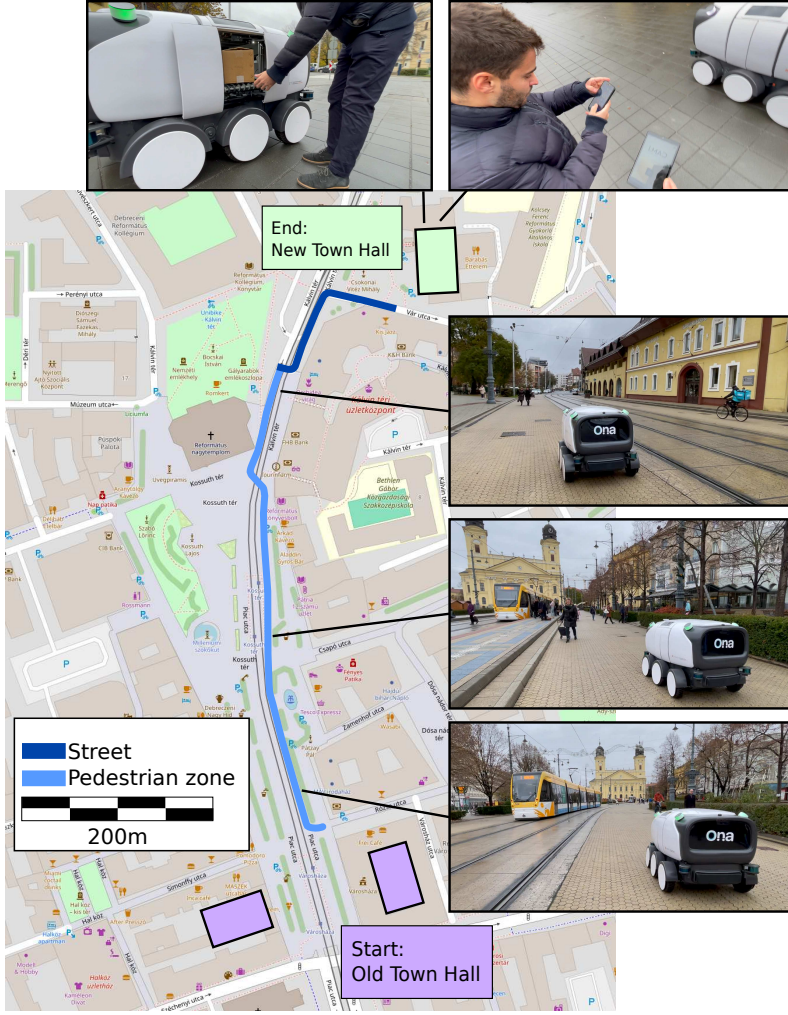
<sup>4</sup><https://serve.postmates.com>

<sup>5</sup><https://otsaw.com/camello>

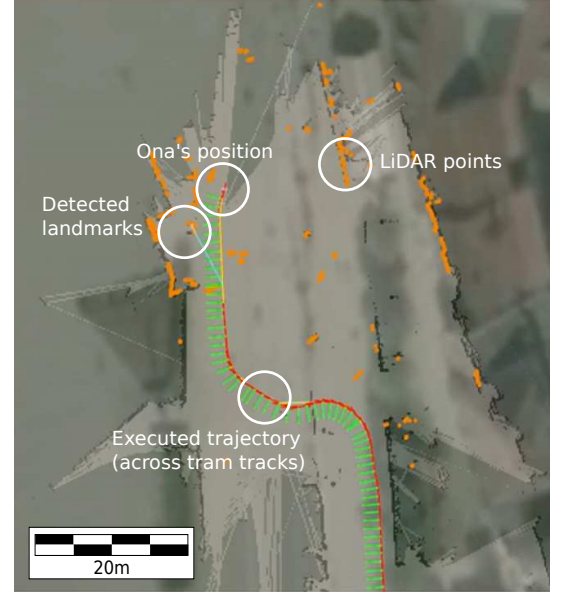
<sup>6</sup><https://bit.ly/roxo-fedex>

<sup>7</sup><https://clevon.com/clevon1/>

<sup>8</sup><https://nuro.ai>



(a) Pilot site and mission details.



(b) Perception of Ona, on top of its map, during the navigation across the tram tracks while sorting Vulnerable Road Users.

Fig. 1: Pilot test at Debrecen (Hungary) in November 2022, using Ona, our autonomous last-mile delivery device.

platform. In contrast to Clevon 1, Ona has been designed for its navigation in pedestrian areas, including for example top bright signaling coloured lights or a large frontal screen for human-robot interaction. With respect to Nuro, Ona is also bit smaller. By design, both Clevon 1 and Nuro favor the street navigation, with Ona meant to navigate in both street and pedestrian zones; with a different navigation stack; and a novel storage management concept. A side from the technical details, Table I provides an overview of the capabilities related to the last-mile delivery challenges, related to the storage; navigation area types; main architectural barriers (e.g., steps or stairs); HRI; and motion models. Once again, Ona is close to the Nuro prototype although Ona has two important extra features: the ability to navigate in pedestrian areas and a motion train based on an Ackerman model (*Ack.*) together with a differential drive (*D.D.*), extending Ona's maneuverability.

Despite all these efforts in terms of technological developments, prototyping and algorithms [6], [7], the AutoNav solutions regarded as state-of-the-art are still far from the maturity needed to safely operate in complex and highly

uncertain intracity scenarios. With the aim at pushing towards a real deployment of our robot Ona, in this paper we present and analyse the following key aspects:

- **Specialized Hardware** able to maneuver and drive at required speeds in both street and pedestrian areas.
- **Localization and navigation** using common sensors for urban delivery operations.
- Usage of virtual tools (i.e., **simulation tools**).
- **Validation in realistic scenarios.**
- **Study of technology acceptance.**

In particular, the rest of the paper is organized as follows. Section II presents our autonomous delivery device, Ona. With a particular emphasis on its design (e.g., the platform or its parcel locker) and localization and navigation software stack. The validation experiments and realistic scenarios are described in Section IV. The technology acceptance (i.e., feedback from bystanders during the experiments) is presented in Section V. Finally, lessons learned and final remarks are drawn in Section VI.



Fig. 2: Examples of last-mile delivery robots under development

	Kiwibot	Starship	Scout	Serve	Camello	Roxo	Clevon 1	Ona	Nuro
<b>Pkg. cap.</b>	↘	↘	↘	↘	→	→	↗	↗	↗
<b>Road nav.</b>	✗	✗	✗	✗	✗	✗	✓	✓	✓
<b>Ped. zone</b>	✓	✓	✓	✓	✓	✓	✓	✓	✗
<b>Step size</b>	< 10cm	< 10cm	< 10cm	< 10cm	< 10cm	< 10cm	< 20cm	< 20cm	< 20cm
<b>Stairs</b>	✗	✗	✗	✗	✗	✓	✗	✗	✗
<b>HRI</b>	↗	→	→	→	↘	↘	↘	↗	↗
<b>Kinematics</b>	Ack.	D.D.	D.D.	D.D.	D.D.	D.D.	Ack.	Ack. D.D.	Ack.
<b>Bi-dir. nav.</b>	✗	✗	✗	✗	✓	✗	✗	✓	✓
<b>LiDAR</b>	✓	✗	✓	✓	✓	✓	✓	✓	✓

TABLE I: Main capabilities of the most famous prototype robots for urban last mile delivery, being "Pkg. cap." the storage capacity; "Road nav." is marked with a green check if the robot is enabled for street driving; "Ped. zone" specifies if the robot is suited for accessing pedestrian-only areas, in contrast to road navigation; "Step size" shows (approximately) the height of a potential obstacle that a robot could be overcome by driving over it; "Stairs" refers to the ability of climbing several consecutive steps; "HRI" evaluates the interfaces of the robot that can enable a correct Human-Robot Interaction (e.g., screen monitors); in "Kinematics" we show the motion model of the robot; "Bi-dir. nav." refers to the capacity of the robot of driving either forward or backwards, being related to the symmetry of the Hardware design; Finally, The "LiDAR" field is to expose the usage of 3D LiDARs in their robotic solutions.

## II. ONA: OUR AUTONOMOUS LAST-MILE DELIVERY DEVICE

Ona has a mobile platform featuring six-wheel drive with steering and Ackerman drive. In particular, it weighs around 200kg and is roughly  $1,8 \times 1,1 \times 1m$  with the outer shell. Ona is an all-electric vehicle and has an autonomy of more than 5h of continuous operation between charges, an autonomy more than enough for last-mile deliveries given the storage capacity shown hereafter. Ona has 4 front traction-only wheels and two rear traction and steering wheels. This type of architecture implies that some of the wheels will skid during its turns. Further, Ona has two (front and back) 22" and a lateral 7" screens, which are meant to interact with a potential customer or pedestrians along the route.

Ona is equipped with a multitude of sensors, each type based on different physical phenomena to provide robustness, in particular:

- **Wheel encoders.**
- **Inertial Measurement Unit (IMU)** providing an estimate of its angular velocity and linear acceleration.
- **Global Navigation Satellite System (GNSS)** to provide accurate global positioning. Even though Ona's GNSS receiver can work in Real-Time Kinematic (RTK) mode together with a base station to improve accuracy, it is currently working in standalone mode.
- **3D LiDARs:** Two 3D lasers of 16 beams are installed in the opposite corners of the robot (front-right and back-left corners). The information from these sensors is used to improve the odometry estimation and create a local map (i.e., point cloud). Section II-A details how this sensor works and how Ona uses it to obtain a robust localization.
- **Depth Cameras:** There are three depth cameras on the front and sides, facing down. The field-of-view of these cameras include the immediate ground next to the robot (in the front and to the sides) and their usage is focused on high-resolution local traversability analysis. These operations run at the camera frame rates (individually processed) and we take advantage of our previous work [8]. To detect potential obstacles, we extract the dense pointclouds from the ground surrounding the robot





Fig. 3: Ona's storage management system with four internal compartments.

and filter the height inconsistencies by using an average filter (performing an integration approximation by observing each pixel in the point cloud at consecutive time intervals and keeping the  $z$  distance average). Then, we down-sample the resulting pointclouds and remove sparse outliers based on the computation of the distribution of point neighbor distances. Finally, we estimate the normal orientation of local planar patches at each point to classify between traversable or not traversable ground. Further, the area in front of the robot is divided into three regions, and the amount of data in each region is analyzed to evaluate the presence of a hole (negative obstacles). Notice that the average frame rate of this depth cameras is 30 frames per second, thus this analysis is mainly used during local maneuvering (e.g., Figure 5b). We refer the reader to [8] for further details. Additionally, a depth-color camera is placed in the front of Ona pointing forward. Here, after thorough outdoor experiments in urban settings, we decided not to use the depth information of the camera due to poor robustness and this camera is mainly used to aid in the teleoperation of Ona (notice how the navigation stack mainly relies on LiDAR readings).

- **Safety 2D LiDAR** used for hardware safety purposes.
- **Bumpers** as a hardware safety feature.
- **Sonars** placed in the front of the wheels and used to detect potential holes or obstacles next them.

Aside from these sensors Ona has several emergency stop implementations (ranging from Category 0 to 2). Further, Ona is equipped with a novel pickup and delivery system to handle the storage, composed of **i)** automatic lateral doors, which will open/close upon request of the mission manager when the pickup/delivery and the costumer are ready; **ii)** automatic package manipulation system, with a belt that will move the packages in and out of Ona; and a 2-axis manipulation device to place or pick the packages from the storage parcels; and **iii)** four storage parcels, designed to retain the packages during navigation phases. All these elements can be seen in Figure 3.

To actuate and interact with the Hardware mentioned above we take advantage of the Robot Operating System<sup>9</sup> (ROS) middleware, which handles all internal communications. Further, we use Behavior Trees [9] as an internal robot task planner, managing the relationships between the mission management; remote control; parcel locker; navigation software stack; and the platform (see the action manager block in Figure 4a). Externally, Ona needs Internet connection to communicate

with the control center, to coordinate all logistic tasks (mission manager) and also to interface with a remote human operator to monitor and assist it (monitoring and teleoperation) when it may get lost or stuck (see the blue arrow connections in Figure 4a). The Ona solution consists of a communications' backpack, a commercially available device with parallel SIM cards (provided by a telecommunications service provider), supporting both 4G and 5G and automatic roaming in almost 200 countries around the world, including Europe.

Regarding Ona's integration in the urban setting, in the last-mile distribution using robots, there exist different categories of human roles and Human-robot Interaction (HRI) tasks, investigated in our conference publication [10] and summarized with: The supervisors, which organize the logistics plans and tasks, optimizing the routes, the number of robots and the delivery of the goods; an operator, a skilled agent that manages the alarms and solves the navigation or deliveries in difficult cases (e.g., teleoperation of the robot from a remote site, in case it is needed); a mechanic/technician, in charge of the maintenance and repair tasks; a peer teammate, who could do cooperative tasks as for example, loading the robot, guiding and accompanying the robot, doing handover tasks or recovering the robot in case of a problem; a peer end-user, which is the costumer that receives the goods; and finally, Bystanders. The bystander role is assumed by the citizens that co-exists in the same environment of the robot and they can fall in different categories: citizens in good health condition; vulnerable citizens as elderly, kids and disabled people; urban services' workers; etc. The robot has to be aware of the type of citizens that could meet and also of pets, bikes, etc. It could also interact with people following the social norms, be aware of them and use communication methods and signs that people can understand. This HRI communication is done by Ona through screen monitors to show messages and drawings (e.g., icons or emojis) and using coloured lights, following a standard coding with green, orange, and red to show the severity of the navigation status as nominal, warning or fatal, respectively. The actual pick-up/delivery system connecting the logistics operator and the peer end-user is based on a set of messages send to the latter with specific codes to validate its authenticity. The HRI interface of Ona is currently minimal for a nominal operation and this topic will be further investigated in the frame of the BotNet project<sup>10</sup> (23S06128-00), financed by the Barcelona city council and La Caixa foundation, which started in November 2023.

<sup>9</sup><https://www.ros.org>

<sup>10</sup><https://www.iri.upc.edu/project/show/317>

## A. Navigation stack

1) *Simultaneous localization and mapping*: To navigate autonomously it is crucial to estimate the ego-motion of the robot and to create an accurate model of its surroundings by analyzing sensor observations [11]. These estimations are handled by the SLAM module, shown in Figure 4a, and detailed in Figure 4b.

Each sensor has its own particularities, and a common visual sensor used in autonomous driving is the LiDAR (light detection and ranging), which has a wide field of view but usually provides non-dense information, requiring the concatenation of readings using the motion estimation to represent the surrounding scenario. LiDAR measurements come at low frame rate, hence its observations are usually fused with those from an inertial measurement unit (IMU). In this sense, Ona fuses the IMU observations with those from the LiDARs with our own sensor fusion architecture, which draws inspiration from LIO-SAM [12] but includes an extra observability module as in our previous work [13]. The input of this module, named LIO-SAM obs. in Figure 4a leverages the LiDAR pointclouds resulting from the traversability analysis (described hereafter) to fuse them with IMU measurements. The output of this initial LiDAR-IMU fusion is then combined with the wheel odometry using a well-known Extended Kalman Filter (EKF) (see the odometry estimation module in Figure 4a).

Even though this odometry estimation was providing fair ego-motion values, it is not globally consistent and subject to localization drift. As a common practice in robotics, a localization outer loop was set to **i)** localize with global consistency Ona (i.e., fusing the odometry estimation with reading of the GNSS module); and **ii)** obtain a local representation of the surroundings (i.e., a map) combining the estimated motion and the laser scans, obtaining a point cloud, to evaluate the required navigation plan and compute control commands. The approach used to compute the global localization and the local map was Cartographer [14], a system able to handle real-time simultaneous localization and mapping (SLAM). In order to give the map Universal Transverse Mercator (UTM) consistency, we added another EKF fusing the readings of the GNSS and the IMU heading estimation, namely UTM calibration module. The connections of this module and the outer global localization loop using Cartographer are shown in Figure 4b.

2) *Traversability analysis*: Ground segmentation is a crucial task for autonomous robots because the quality of its results conditions all the subsequent higher-level processing stages like obstacle detection, path planning, or localization, to name a few. Here, we take advantage of our probabilistic graph-based real-time ground segmentation algorithm [15] that, taking as input an instantaneous 3D point cloud, generates a probabilistic model representing the ground surface and its traversability in real-time. The algorithm builds a graph while exploring the point cloud from its origin –where a solid prior of the ground level is available– to its limits. It does so by fitting small planes to the data in a probabilistic filter fashion: extract local data, use the prior to reject outliers, update the plane estimation with inliers and explore new

planes propagating the posterior as a prior for the new nodes. This approach permits to segment the point cloud taking into account the uncertainties of the model explicitly. Moreover, during the graph building stage, we extract some statistic descriptors –including intensity information– to refine the traversability analysis by using a Random Forest model ([16]) to distinguish different types of surfaces like roads, sidewalks, low vegetation areas, etc. For the sake of simplicity, we refer the reader to [15] for further details on this algorithm and its thorough evaluation. An example of this algorithm processing the ground segmentation in The Barcelona Robot Lab<sup>11</sup> (North campus of the UPC) is shown in Figure 5a, where ground and obstacle points are depicted in blue and green, respectively. Vertices and edges of the roadmap are depicted in brown. The traversability analysis is leveraged by the planning modules (for both pedestrian and road navigation, shown in Figure 4a) and the segmented pointcloud including the static object above the ground is fed into the LiDAR-IMU fusion algorithm (see Figure 4b).

3) *Global map handling and path planning*: The mission manager, handling the business operation, provides Ona with the global UTM coordinates of the pickup and delivery points. Between these points, Ona extracts a global path by consulting a map service and running path planning algorithms [17]. Hence, with the robot state expressed with respect to a global reference frame, thanks to the global localization approach mentioned above, Ona needs to access a global map provider to compute the navigation trajectories (see the online urban map provider Figure 4a).

In particular, to compute the global path, we take advantage of the Open Street Map (OSM) service<sup>12</sup>. Ona has a local copy of the OSM map of the region where the mission is expected to happen. Notice that OSM is an open-source online map resource that can be edited by any user. Hence, to avoid incomplete or wrong data, we use this local copy of the OSM zone where Ona will navigate, allowing us to manually check the correctness of its information (or even complete it) before its usage. We consider this a bearable task as it is just required once in each area. With this OSM map, Ona has a ROS node to extract the routes from one point to another. The output of this node allows Ona to see if the route has pedestrian or road segment paths (a classification already existing in OSM maps), thus it can split the required route into parts (split path) to define different sequential trajectory goals to navigate considering the specifics of one type or the other (road or pedestrian areas), thus Ona can then select a path planner specialized for road or pedestrian area navigation [18]. This specialization arises from the fact that the pedestrian areas are less structured than conventional roads and, also, to consider potential vulnerable road users while computing the global path. Once the segments have been divided into road and pedestrian ones (shown in Figure 4a), we execute different planners to compute a global trajectory for the robot, which are briefly described in the following:

- **Road Navigation**: It plans a shortest path/lane from a

<sup>11</sup><http://www.iri.upc.edu/research/webprojects/pau/datasets/BRL>

<sup>12</sup><https://openstreetmap.org>

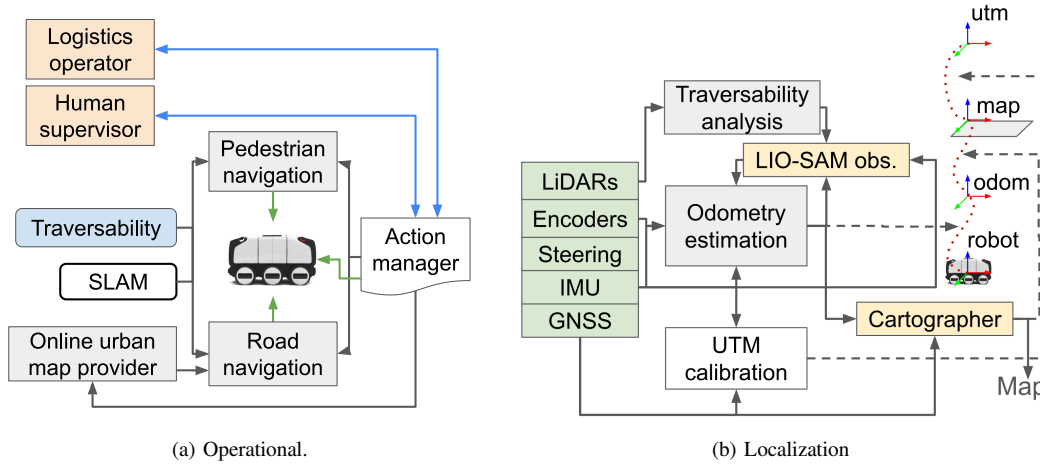


Fig. 4: Main Software modules of Ona's autonomous navigation pipeline

current location to the final parking spot following traffic rules define in the map. The path is given as an array of the lane id from the starting to the goal position. It is to note that the global path is the center lane line and include the path from/to the parking spot through the drivable area defined in the map.

- **Pedestrian area navigation:** This planner will generate a path from the robot's current position to a desired goal pose. Paths are generated by combining a series of "motion primitives" which are short and kinematically feasible motions. Figure 5b-right shows an example of these motion primitives. The plans result in smooth paths that consider the robot orientation. One of the main features of this planner is that it generates a set of maneuvers when a simple path cannot be found, as shown in Figure 5b-left.

The global path is only computed once for every mission and, to follow it, Ona has a local planner with reactive behaviors, especially suited for vulnerable road users (VRU). This global- local-planner relationship follows the ROS move-base architecture <sup>13</sup>.

The local planner, in contrast to the global one, re-plans several times per second in order to adapt to unexpected situations, for instance in the presence of VRU. From the traversability analysis, Ona generates global and local cost maps as occupancy grids, where local planner will search for a suitable local plan in every control cycle. The Ona motion model used in the local planner is an Ackerman design with a custom modification to consider Spline trajectories to better smooth the paths and adapt to the real Ona motion. A part from dealing with obstacles and VRU, these planners have recovery behaviors defined to recover in the event of an unfeasible plan. These behaviors are mainly related to re-planning tasks and, in the event that no valid plan is found, the task manager will trigger a flag to require assistance from the operator using the Behavior Trees' machinery [9].

4) *Remote monitoring and teleoperation:* Ona is able to communicate bidirectionally to an external control center sup-

porting ROS2 (Figure 4a). In particular, this communication has been designed to minimize the effects of the following constraints:

- **Bandwidth (BW):** We compress the data from the robot to upload a maximum of  $1\text{Mbit/s}$  (the upload BW is usually the most constrained by the telecommunications companies), consisting of robot state messages and the front image for visual feedback purposes. The data downloaded in the robot has a rate of  $8\text{Mbit/s}$  and consist of control commands. These rates can be accomplished tuning the ROS2 QoS network controller.
- **Latency:** To teleoperate Ona the latency should be (and is) less than  $300\text{ms}$ . This latency parameter can be tuned in the ROS2 throughput controller.

Notice that even though 4G is already present in all major and medium cities of western countries, the usage of 5G signaling might not be a possibility nor distributed uniformly. Hence, an initial verification that the 4G network has no substantial shadows is required when deploying Ona in a new city.

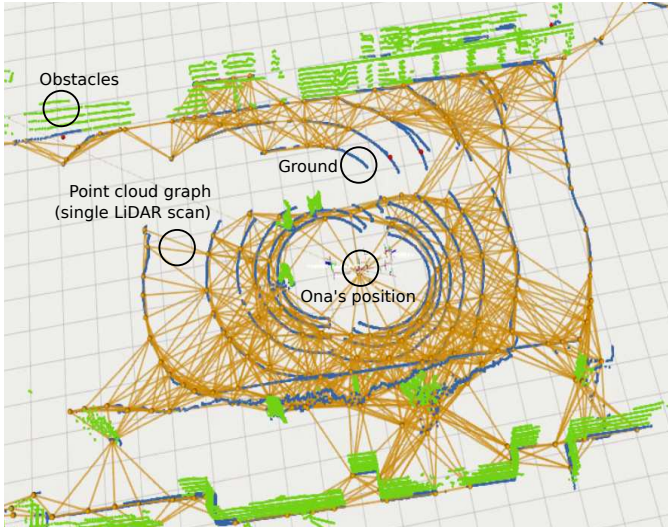
5) *Logistic's mission manager:* Ona can easily be integrated in any major logistics's operator Software. Currently, Ona can send to the mission (business) manager any information related to its state, the state of the parcels or the delivery/pickup status. In contrast to other existing last-mile robots, Ona has a unique compartment mechanism that fetches parcels and opens/closes doors automatically, thus the exchange of information between the mission manager and Ona can be reduced to the initial and final locations (UTM coordinates) and the pickup/delivery synchronization (i.e., operation trigger). Ona can also send live navigation data to the mission manager for tracking purposes.

### III. POTENTIAL USE CASES

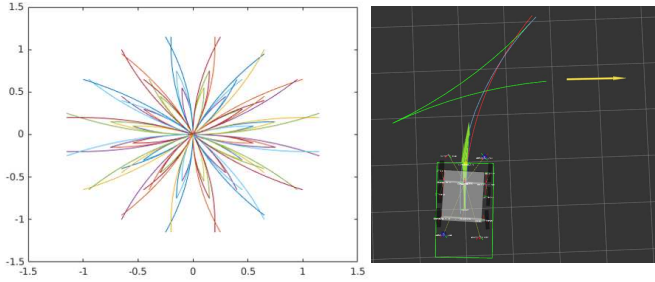
Ona has been designed with a set of initial use cases to guide its developments, in particular:

- **Follow-me mode:** Ona would follow a person or another vehicle at a given distance, increasing the weight that could normally be carried (i.e., Ona autonomously carrying heavy loads).

<sup>13</sup>[http://wiki.ros.org/move\\_base](http://wiki.ros.org/move_base)



(a) Output sample of the ground segmentation algorithm used for traversability analysis, processing the merged instantaneous LiDAR readings of Ona, during an experiment in The Barcelona Robot Lab (North campus of the Universitat Politècnica de Catalunya).



(b) Example of motion primitives (left) and a generated path (right) computed by the pedestrian area global planner (green lines: global path; blue line: current segment executed by the local planner; red line: actual motion command generated by the local planner).

Fig. 5: Perception and planning details of Ona.

- **E-commerce parcel delivery and pick-up:** Ona receives parcels (pick up) and then finish the delivery (navigation) to an end customer (delivery).
- **Automated food deliveries from supermarkets/urban markets/restaurants:** Service to collect groceries from supermarkets in a city (i.e., employees of the markets could feed the Ona parcels with customer orders) and to deliver them to the end customers.
- **Autonomous vending machine (e.g., food, kiosk):** Ona would be loaded with groceries or other types of products and travels throughout the streets of a given service region. If a customer wants to buy something, he/she has to hail the robot.
- **Logistics on industrial site:** Similar to the above use cases but, instead of operating in open and public environments, Ona usage could be restrained to closed private industrial sites.
- **City mapping and cadastration:** As is constantly receiving information from its surroundings and using its local representation (i.e., a 3D map). Then, if adequately treated, this data could provide useful insight to cities

(potholes in roadway, obstacles on sidewalks, etc.).

- **Autonomous transport of recyclable material:** Automated garbage transportation with Ona to make waste management more flexible and efficient.
- **Passenger transport:** Ona base platform could be used (with some modifications) as an autonomous individual pod for passenger transport.
- **Information point for tourists or tourist guide:** Information about the city and its monuments could be displayed in Ona's screens and shown to tourists, either statically or combined with the follow-me mode to act as a tourist guide.

All these use cases take advantage of Ona's capabilities related to the autonomous navigation, its capacity to transport heavy packages or to interact with surrounding pedestrians. During the first tests (see the following section), we focused on the E-commerce parcel delivery and pick-up use case. This operation purpose is the most important one for the first implementation of Ona in real life and, even though the Ona capabilities will be validated with this E-commerce use case, they can be easily exported to most of the other use cases with minimal adaptations.

## IV. VALIDATION AND EXPERIMENTS

### A. Simulation tools

We developed two different simulations tools for Ona. On the one side, we have a simulated environment based on the Gazebo<sup>14</sup> simulator, already integrated with ROS. On the other side, we created a software layer to interact with the PTV Vissim traffic simulator<sup>15</sup> to test navigation conditions and interactions with the environment.

Regarding the PTV Vissim simulation, the navigation Software of Ona can virtually "move" the simulated Ona model in the PTV Vissim virtual environment and react accordingly to quickly and cost-effectively catch bugs and improve the code while testing Ona in a traffic-based simulation, which allows to simulate specific predefined urban areas and realities (see for instance the digital twin of the pilot site in Esplugues de Llobregat, shown in Figure 6 - top), easing the post analysis and evaluation, for example, with more/less obstacles, more/less crowded, with/without traffic/bikes/scooters, with unexpected events, etc. An example of the semantic information that can be used in the PTV Vissim simulated environment is shown in Figure 6 - bottom. The usage of these simulations are very useful to consider other vehicles and pedestrian flows while solving their interactions within the shared spaces or conflict areas.

### B. Real tests in representative urban scenarios

Moving autonomously in an urban setting presents several challenges and open problems. These are complex 3D scenarios and the space has to be safely shared with bystanders, animals and other vehicles [19]. Among all challenges, we find of paramount importance the following:

<sup>14</sup><https://gazebo.org>

<sup>15</sup><https://www.myptv.com/en/mobility-software/ptv-vissim>



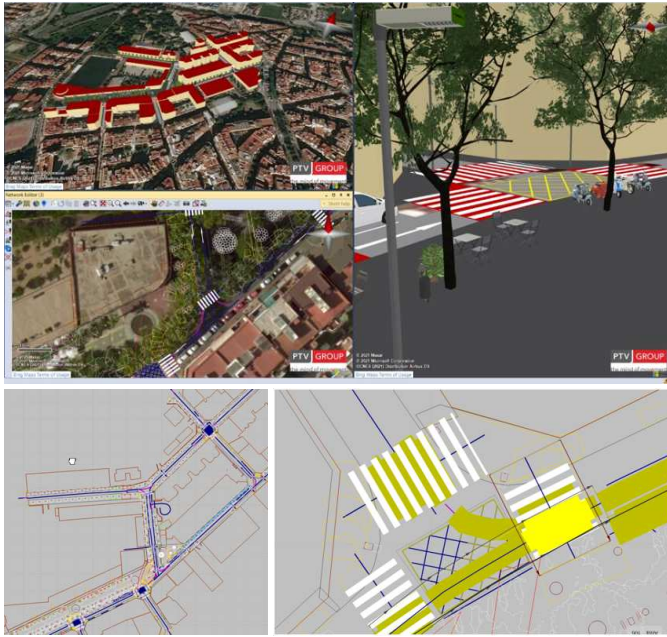


Fig. 6: Virtual reality environment (top) and Semantic information (bottom) codified in the virtual urban environment of the Esplugues de Llobregat pilot site (Spain) using PTV Vissim.

- Presence of dynamic objects, people, animals or vehicles:** As a dynamic vehicle, Ona navigates over a static urban setting shared with others that may also have their own dynamic motion. In this case, we can remove them using [20] from the data incoming to the SLAM approach that estimates the vehicle's location. By doing so, the respective observation models are simpler, faster and can be processed as static scene landmarks. On the other hand, the segmented dynamic elements can be added to an Occupancy Grid Map (OGM) of the navigation stack, which is used by the planning modules to navigate safely and evade obstacles. Drawing inspiration from [21], this trajectory estimation of dynamic elements can be included in the OGM to better optimize the maneuvering path.
- Weather conditions:** All robots perceive and represent their surroundings using sensors based on different physical phenomena. Such events can be affected by weather conditions including, for instance, sunny or rainy days; or the presence of fog. As time goes by, the scene might even differ between seasons. The robustness on processing such sensory data disturbed by the weather conditions is still an open challenge in robotics. Although we tested Ona in two different cities (Esplugues de Llobregat, Spain, and Debrecen, Hungary) Ona behaviours have shown no operational difference as our focus has been in navigating in nominal weather conditions, during daylight without rain, nor snow or fog.
- Existence of architectonic barriers:** Even-though urban scenarios are quite structured, the presence of narrow corridors, steps, uneven floors or stairs are common. This not only hindrances the people's right to mobility (e.g., when using a wheelchair) but also affects the navigation

of vehicles. Overcoming such obstacles commonly entails specific robotic designs (e.g., legged robots) and present a compromise between accessibility and mobility and strongly depends on the vehicle's Hardware configuration.

- People's curiosity:** A robot navigating among human beings attracts people's curiosity. Robotics is a relatively new discipline compared to the western cultural developments and bystanders cannot have a clear interaction model while being next to Ona in the street. Such lack of prediction, together with inherent human's curiosity produces an attraction effect on people, which usually surround the robot and automatically blind its perception. Even while Ona is moving, the public follows it, increasing the complexity of the localization approach as the ratio between dynamic and static obstacles is not in favor of solving properly the SLAM approach. Ona's current solution relies on stopping, advertising the situation to bystanders by using coloured lighting and then resuming the mission.

Overall, the development of Ona has been focused on solving its navigation tasks in urban nominal conditions (e.g., with a regular pavement or lane, and without big steps in the way of the robot), leaving as future work the study of edge cases, the effect of strong season differences or unpredicted situations (e.g., sudden road blockades). In all these situations, Ona can cancel its navigation status if the path is unfeasible; if the localization is unreliable; or if any other issue arises, asking for support and waiting for remote operation to overcome the main issue before resuming its mission.


Several integration and testing efforts took place in different locations, chosen aligned with the development and testing phases of Ona (reflected with dates):

- Barcelona Robot Lab (Spain), January to March 2022:** The Barcelona Robot Lab encompasses an outdoor pedestrian area of 10.000 m<sup>2</sup>, in the North Campus of the Universitat Politècnica de Catalunya. The area has moderate vegetation and intense cast shadows, challenging conditions for the computer vision algorithms.
- Esplugues de Llobregat city center (Spain), May and June 2022:** Located in the busy Esplugues de Llobregat, with one of the highest population density in Europe.
- Debrecen city center (Hungary), October and November 2022:** Located in Debrecen, this is the second place for open-public tests.

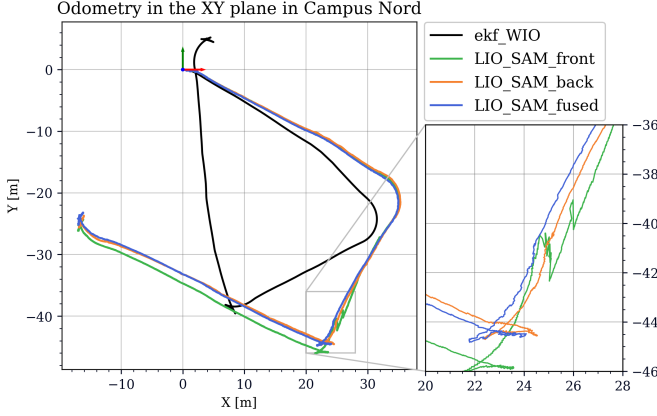
It is worth mentioning that in all these scenarios gathering navigation (e.g., localization) ground-truth data was not possible. For instance, no GNSS (Global Navigation Satellite System) signal was robust enough due to the nature of the scenes. In the Barcelona Robot Lab and the dense Esplugues de Llobregat (both in Spain), the GNSS had huge multi-path interference and aliasing between tall buildings. In Debrecen existed a poor signal reception, with a Virtual Reference Station (VRS) being created in Budapest (200 kms away from the testing site), disabling our Real-Time Kinematic GNSS (RTK) solution as a ground-truth of the global localization system.

Given the importance of the localization module, whose

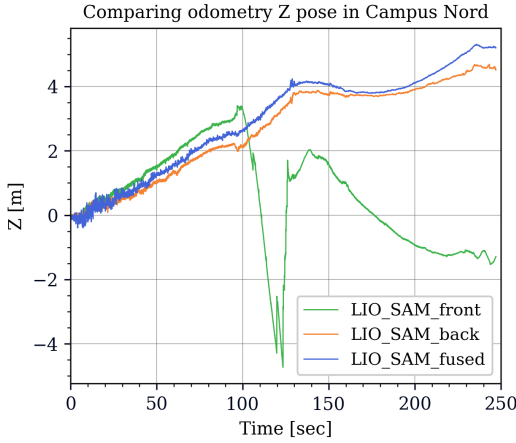




(a) Barcelona Robot Lab test site.



(b) Odometry estimations from three instances of "LIO-SAM obs." and an EKF-based Wheel-Inertial approach (WIO).



(c) Odometry results in z.

Fig. 7: Barcelona Robot Lab experiments, Universitat Politècnica de Catalunya, Spain.

errors can cascade down the rest of the navigation pipeline (e.g., planning or control modules), and the lack of ground-truth data in the real scenarios, we focus the following in presenting the localization results of the "LIO-SAM obs. module", being the most crucial approach within the SLAM scheme shown in Figure 4b. The evaluations hereafter include comparisons of using the localization system with different on-board sensor means or algorithms (e.g., EKF-based odometry or LIO-SAM). In all these experiments, the actual starting and end poses of the robot were equal, providing the drift error of our approach.

1) *Barcelona Robot Lab*: The first experiments with Ona were conducted around the Barcelona Robot Lab (BRL), a controlled site, shown in Figure 7a, where different robotic hardware configurations could be tested. The specific features and challenges of this site are shown in Table II.

Among other tests, in BRL we evaluated the differences between running LIO-SAM [12] (see Section II-A) with the front LiDAR sensor (*LIO\_SAM\_front*); the back LiDAR sensor (*LIO\_SAM\_back*); or with the merged point-cloud from both sensors (*LIO\_SAM\_fused*); compared to a traditional fusion of the Wheel Odometry and IMU readings using an EKF (*ekf\_WIO*). These comparisons pretend to assess the best strategy to estimate the ego-motion (i.e., odometry estimation) with the available Ona means and considering the trade-off between computation and precision (e.g., *ekf\_WIO* is lighter than processing a merged pointcloud in *LIO\_SAM\_fused*).

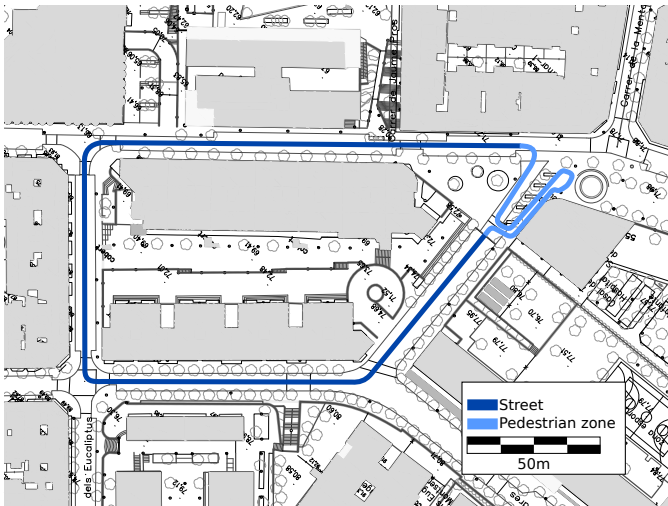
Figure 7b plots the xy-odometry of the three LIO-SAM instances, alongside the *ekf\_WIO* for reference. The zoomed region highlights the moment the robot is at the end of a long hallway, which has a slight upward incline. As a consequence, the front 3D LiDAR was only able to pick up the vertical walls. The degradation in the environment is responsible of the zig-zag behavior in the *LIO\_SAM\_front* odometry, as well as the jumps in the z-estimate seen in Figure 7c. The back LiDAR could still observe corners and columns outside of the hallway most of the way through. These features resulted in a better xy-odometry estimation. In addition, the upwards slope meant that the back LiDAR could see the floor better than the front LiDAR, thus avoiding the large changes in z-estimate as seen with the front LiDAR LIO. By merging the point-clouds, the robot can observe  $360^\circ$  surrounding it. And, as a result, the state estimation suffers less from the LiDAR degraded environment compared to using a single LiDAR, hence the extra computational burden is acceptable not to jeopardise a mission.

2) *Esplugues de Llobregat city center*: This test site covers a 750m track that mixes roads and pedestrian areas. The chosen navigation route is depicted in the map shown in Figure 8a. The initial tests in this site included open street navigation in the streets marked with dark-blue in Figure 8a. As the previous tests in BRL already emulated a pedestrian area, our focus here was on street navigation.


In this scenario, the multitude of static features made the scene very observable for the LiDAR. As a result, the total localization error after 500m was below 30cm (Figure 8b-left). In this case, first, we compared the effect of enabling or disabling the loop closure factor in LIO-SAM. In theory, enabling the loop closure should decrease the final error, as the loop closure serves to eliminate drift accumulated when the same features are observed again. A disadvantage with loop closures is that we can no longer consider it an odometry estimation source and it loses its estimation smoothness. As shown in Figure 8b, enabling or disabling the loop closure factor appears indistinguishable in the xy odometry. The final error in the xy plane for both cases is less than 30cm. Secondly, we tested here a new fusion solution, WILO, as the results of a local EKF that fuses wheel, inertial, and laser odometry. The laser odometry for WILO will be from the LIO-SAM instance

	Barcelona Robot Lab	Esplugues de Llobregat city center (Spain)	Debrecen city center (Hungary)
<b>Pavement layout</b>	<i>Several levels and underpasses</i>	<i>Road and pedestrian areas</i>	<i>Mainly pedestrian areas</i>
<b>Vegetation</b>	<i>Mild</i>	<i>Moderate</i>	<i>Moderate</i>
<b>Architectonic barriers</b>	<i>Mild</i>	<i>Moderate</i>	<i>Severe (e.g., non-segregated tram)</i>
<b>GNSS coverage</b>	<i>Intermittent</i>	<i>Intermittent and not reliable (aliasing)</i>	<i>Lack of precision (RTK VRS 200kms away)</i>
<b>Sunlight exposure</b>	<i>Severe, existence of cast shadows</i>	<i>Severe dynamic range (direct sunlight and dark shadows between tall buildings)</i>	<i>Mild (cloudy)</i>
<b>Scene dynamics</b>	<i>Moderate, few bystanders</i>	<i>Severe, lots of bystanders</i>	<i>Severe, lots of bystanders, bicycle riders and shared tram route</i>

TABLE II: Main perception and navigation challenges of the test scenarios.



(a) Map of the test site in the center of Esplugues de Llobregat (dark blue: road; light blue: pedestrian zone).



(b) XY localization for different odometry estimators.

Fig. 8: Esplugues de Llobregat experiments, Spain.

with the loop closure factor disabled. In addition, the wheel-inertial odometry (WIO) showcases, once again, the necessity of incorporating LiDAR odometry to the localization problem.

3) *Debrecen city center*: In November 2022 we did a second pilot demonstration in the city center of Debrecen,

Hungary. Here, the length of the pilot route was approximately 600 meters and it covered one of the most frequented areas in Debrecen. Most parts of the route were completely pedestrianized, with car traffic only affecting the last 150m. One of the improvements with respect to Esplugues was on the consistency between the global (OSM with UTM coordinates) and local (SLAM) maps. For this, we included observations of external geolocalized landmarks to reduce the state estimation uncertainty in the SLAM solution. The landmarks at the starting location of Ona were of special importance (notice how we could normally add them in a robot hub were several Onas could be charging and waiting for a last-mile delivery mission). The main objective of this pilot test was to showcase a full delivery mission, delivering an actual package. The route and some of the mission stages are shown in Figure 1a.

Overall, the Debrecen tests and demonstration were a success. Ona performed as expected conducting the last-mile delivery mission. Figure 1b shows an example of the Ona perception system while sorting VRU in the middle of the delivery route. Notice the VRU as orange dots in this Figure 1b, next to the Ona trajectory (depicted as a yellow line with a sequence of 2D coordinate frames, with red x-forward and green y-left axis) and how Ona modified its route to keep a safe distance from them while moving forward. This exact case happened in the middle of the route when a tram vehicle did bring several pedestrians to a tram station just in front of Ona. See in Figures 1 and 1a how Ona is sharing the route with the city tram and several pedestrians. The satisfaction of VRU was already assessed in the Esplugues pilot test thorough surveys, as described in the following section.

## V. TECHNOLOGY ACCEPTANCE

During the tests in the city center of Esplugues de Llobregat, we distributed a survey to the bystanders to assess the technology and human acceptance of Ona. The details of such technology acceptance can be seen in [10]. This survey was distributed in Catalan and Spanish print version to the bystanders. The English version was made for the online

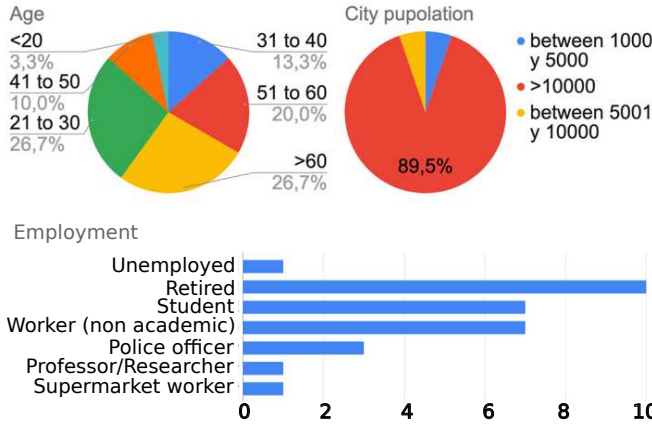


Fig. 9: Classification of bystanders that fulfilled the technology acceptance survey during the pilot test in Esplugues de Llobregat.

survey through an online form. Further, we added supporting material to the attending public in the form of videos and images.

The statistics of bystander types are shown in Figure 9 and the complete survey and analysis of each answer can be found in: <https://bit.ly/survey-ona>. As a summary, the analysis lead to the following:

- There is a high acceptance on the use of robots for last-mile deliveries.
- The public suggests a first integration in industrial or university sites.
- These robots might cause some degree of interference in the city daily life; thus, they prefer scheduled daily autonomous deliveries.
- There is a high confidence on the autonomy of the delivery devices and high confidence on the delivery success of robots like Ona.
- A free navigation of the robot is clearly preferred.
- The general public believes the cities require some architectonic modifications to allocate for these types of robots, and they would prefer to integrate the new infrastructure in the current urban furniture.
- Overall, these delivery robots are not seen as annoying (e.g., considering noise or sharing the space) and the public highly values the ecological benefits of the possibility of returning the packaging waste inside the robot right after the delivery.

As general comments, there is a clear acceptance on the usage of these robots although there is also a clear worry about unemployment generated by its integration in our daily lives during the initial phases. Interestingly enough, looking at the conclusions of our previous work [10], senior individuals tend to exhibit reluctance when it comes to the introduction of new technology in the urban public space.

Overall, the acceptance of robotic last-mile distribution is notably high in urban public spaces, particularly in areas with lower population density. However, this existing public space is not fully equipped to seamlessly incorporate this emerging technology and requires adaptation. The survey indicates a

significant willingness to allow the robot's free navigation, as opposed to confining it to a segregated route, as long as safety measures are assured. Essential infrastructure elements such as charging points, cameras, and sensors could be seamlessly integrated into the current street furniture.

On the whole, Ona has not arised a sense of insecurity among pedestrians in open spaces like public squares. However, in more confined areas such as sidewalks, there is a heightened perception of insecurity. It is also important to remark that the tests with Ona have not disrupted the tasks performed by other entities in the public space. The circulation of the distribution robot has not resulted in noise or visual pollution within the public space and the peer teammate is acknowledged not as a separate entity from the robot but as a companion in its vicinity, overseeing multiple units and promoting integration.

## VI. LESSONS LEARNED AND FINAL REMARKS

This paper presents Ona, our last-mile delivery robot. Aside from describing the Hardware and sensors, we provided insights on the navigation software stack, related to SLAM, traversability, planning and global map handling. In contrast to other existing delivery robots, Ona is meant to navigate in both street and pedestrian areas, incorporating a mid-sized storage system that can automatically pickup or deliver packages.

Regarding the SLAM system, we leveraged a state-of-the-art approach (i.e., LIO-SAM [12]) fused with measurements of the wheel odometers (i.e., WILO). This integration has shown crucial, as the WILO final drift was under 30cm after 500m. The reduction in drift in the local odometry means that a fast and accurate GNSS signal becomes less important when navigating, as the global corrections can occur less frequently. This is especially critical in the urban environments where Ona will make deliveries, as the city buildings often block or degrade the GNSS signal.

The use of the Ona simulator resulted crucial in the initial development phases, especially the Gazebo-ROS simulation to test and debug the navigation algorithms. The initial viable tests of Ona navigating in a representative scenario could also be tested thanks to the digital twin of Esplugues de Llobregat programmed in the PTV Vissim software. These virtual tools have resulted very valuable and will also be for future tests. Thanks to their compatibility with the real robot software, these considerably reduce the development time of new capabilities.

We validated Ona in three different scenarios. First, initial tests were conducted at the Barcelona Robot Lab (Universitat Politècnica de Catalunya), to prepare the Ona platform and test the initial sensor configurations and methods. Then, the official Ona presentation was in Esplugues de Llobregat (June 2022). An event attended by Spanish national authorities like the Minister of Transportation. Considering that Esplugues de Llobregat has one of the highest population densities in Europe, the test results were very satisfactory towards the deployment of autonomous last-mile delivery robots. A second pilot demonstration was carried out in Debrecen (Hungary) in November 2022. This time, the city mayor attended the



event and a complete delivery mission was carried out. In this case, Ona shared the route with tram vehicles as well as a considerable number of pedestrians.

As future work, all modules mentioned in this document have room for improvements, especially in the direction of providing robustness. The main expected change in the near future is to improve the mobility of the Ona platform. A second version of the Ona prototype has already been built and the software stack will be migrated shortly. The new Ona not only improves the maneuverability by using a double Ackerman six-wheel system but also has a better locker system. Additionally, it is worth mentioning that we are working with the Spanish Dirección General de Tráfico to obtain a certification for Ona, which is the first vehicle of its kind in Spain and still there is no legislation for this type of robots.

## VII. ACKNOWLEDGEMENTS

This work has been partially supported by the EU LogiSmile project (EIT-UM-2020-22140); by the Spanish Ministry of Science and Innovation under the projects AUDEL (TED2021-131759A-I00, funded by MCIN/AEI/10.13039/501100011033 and by the "European Union NextGenerationEU/PRTR") and LENA (PID2022-142039NA-I00, funded by MCIN/AEI/10.13039/501100011033 and by "ERDF A way of making Europe"); by the BotNet project (23S06128-00), funded by the Ajuntament de Barcelona and Fundació la Caixa; and by the Consolidated Research Group RAIG (2021 SGR 00510) of the Departament de Recerca i Universitats de la Generalitat de Catalunya. We also thank the Spanish Ministerio de Universidades, the European Union-Next Generation and the University of Alicante for the financial support (MARSALAS21-08).

## REFERENCES

- [1] A. Sharma, P. Zanotti, and L. P. Musunur, "Drive through robotics: Robotic automation for last mile distribution of food and essentials during pandemics," *IEEE Access*, vol. 8, pp. 127 190–127 219, 2020.
- [2] M. Viu-Roig and E. J. Alvarez-Palau, "The impact of e-commerce-related last-mile logistics on cities: A systematic literature review," *Sustainability*, vol. 12, no. 16, 2020.
- [3] C. Lemardel, M. Estrada, L. Pagès, and M. Bachofner, "Potentialities of drones and ground autonomous delivery devices for last-mile logistics," *Transportation Research Part E: Logistics and Transportation Review*, vol. 149, p. 102325, 2021.
- [4] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, "A survey of autonomous driving: Common practices and emerging technologies," *IEEE Access*, vol. 8, pp. 58 443–58 469, 2019.
- [5] N. Boysen, S. Fedtke, and S. Schwerdfeger, "Last-mile delivery concepts: a survey from an operational research perspective," *OR Spectrum*, vol. 43, no. 1, pp. 1–58, 2021.
- [6] F. Ingrand and M. Ghallab, "Deliberation for autonomous robots: A survey," *Artificial Intelligence*, vol. 247, pp. 10–44, 2017, special Issue on AI and Robotics.
- [7] S. Fakoorian, A. Santamaria-Navarro, B. T. Lopez, D. Simon, and A.-a. Agha-mohammadi, "Towards robust state estimation by boosting the maximum correntropy criterion kalman filter with adaptive behaviors," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5469–5476, 2021.
- [8] A. Santamaria-Navarro, E. H. T. Avilés, M. Morta, and J. Andrade-Cetto, "Terrain classification in complex three-dimensional outdoor environments," *J. Field Robotics*, vol. 32, no. 1, pp. 42–60, 2015.
- [9] M. Colledanchise and P. Ogren, *Behavior Trees in Robotics and AI: An Introduction*, 1st ed. USA: CRC Press, Inc., 2018.
- [10] A. Puig-Pey, J. L. Zamora, B. Amante, J. Moreno, A. Garrell, A. Grau, Y. Bolea, A. Santamaria-Navarro, and A. Sanfeliu, "Human acceptance in the human-robot interaction scenario for last-mile goods delivery," in *IEEE International Conference on Advanced Robotics and Its Social Impacts*, 2023, pp. 32–39.
- [11] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [12] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, "Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping," *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5135–5142, 2020.
- [13] A. Tagliabue, J. Tordesillas, X. Cai, A. Santamaria-Navarro, J. P. How, L. Carlone, and A.-a. Agha-mohammadi, "Lion: Lidar-inertial observability-aware navigator for vision-denied environments," in *Experimental Robotics*, B. Siciliano, C. Laschi, and O. Khatib, Eds. Cham: Springer International Publishing, 2021, pp. 380–390.
- [14] W. Hess, D. Kohler, H. Rapp, and D. Andor, "Real-time loop closure in 2d lidar slam," *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1271–1278, 2016.
- [15] I. d. Pino, A. Santamaria-Navarro, A. G. Zulueta, F. Torres, and J. Andrade-Cetto, "Probabilistic graph-based real-time ground segmentation for urban robotics," *IEEE Transactions on Intelligent Vehicles*, pp. 1–14, 2024.
- [16] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [17] R. L. Guimarões, A. S. de Oliveira, J. A. Fabro, T. Becker, and V. A. Brenner, *ROS Navigation: Concepts and Tutorial*. Cham: Springer International Publishing, 2016, pp. 121–160.
- [18] G. Ferrer and A. Sanfeliu, "Bayesian human motion intentionality prediction in urban environments," *Pattern Recognition Letters*, vol. 44, pp. 134–140, 2014, pattern Recognition and Crowd Analysis.
- [19] R. Kümmerle, M. Ruhnke, B. Steder, C. Stachniss, and W. Burgard, "Autonomous robot navigation in highly populated pedestrian zones," *Journal of Field Robotics*, vol. 32, no. 4, pp. 565–589, 2015.
- [20] L. Schmid, O. Andersson, A. Sulser, P. Pfrendschuh, and R. Siegwart, "Dynablox: Real-time detection of diverse dynamic objects in complex environments," vol. 8, no. 10, pp. 6259 – 6266, 2023.
- [21] G. Ferrer and A. Sanfeliu, "Proactive kinodynamic planning using the extended social force model and human motion prediction in urban environments," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 1730–1735.