Robots in the Wild
Report II: Service Robots in Public Places

By Pernille Maja Carlsen, Peter Hommel Østerlund and Cathrine Hasse[*1]*

Introduction

This report is the second of five presenting studies of "robots in the wild." Report I problematized autonomy in robots as relational, situated, and often overstated. Report II examines how this plays out when robots meet "newcomers" in public spaces.

As part of the project RoboSAPIENS a group of anthropologists have studied how close the robots of today are to autonomy when studied in real-life settings. The purpose of this study of robots at workplaces, in fields, at sea, and in public places like restaurants and shops is to explore real human-robot interactions as contrasted with the more staged versions found at conferences and in media imaginaries. The findings of the first four reports (Report I-IV) are summed up in Report V, where we present the overall conclusions and discuss what these studies can tell us about a future with RoboSAPIENS technology.

Many of the places, where we have done fieldwork, have robots moving around among people in, for example, airports, educational institutions, shops, museums, hospitals, and restaurants. In this sub-report we look at places that have what we might call public service robots in situations where the robots also meet people who are not used to working with them in public spaces.

In these public spaces we encounter, for example, robot waiters, cleaning robots and robots that transport things back and forth, as well as service robots that can tell you whether you are going the right way. The people in these interactions include those who move daily in the public space – such as airport staff, office workers, and human waiters. But the robots also meet people who may never have seen robots before, such as travellers, restaurant guests, customers in a shop or students. This report covers our visits to several such locations, including restaurants, offices, and shops.

---

[1]Student helpers especially Amalie Rævsbæk Birck and Stephan Holmberg-Hansen participated in gathering the data analysed in these reports.

Public places as a stress test for autonomy

Public spaces are analytically important because they put robots under a particular kind of pressure. Unlike workplaces where staff gradually become "experienced," restaurants, shops, receptions, and museums are environments of constant turnover: most people the robot encounters are newcomers, and interactions are brief. This means trust must be established instantly, and the gap between marketed autonomy and practical dependence becomes visible very quickly. In these settings, small disturbances, as when layout changes, crowds, glass walls, children's movements get in the way, immediately surface for the robots as interruptions, workarounds, or "anomalies." Publicness therefore functions as a worst-case scenario for robot autonomy, not because robots necessarily fail more often, but because the social and spatial conditions expose how much autonomy is produced in practice through human adaptation.

In the first examples we visit restaurants where robots help side by side with human waiters to deliver food to guests.

Restaurant 1 - Waiter Robots

The restaurant – a shared space for robots and humans

We have visited several restaurants, where we have observed robots in practice as well as made interviews with restaurant owners, gusts and waiters.

Restaurant 1, the first restaurant we visit, is a typical medium-sized restaurant with room for up to 100 guests inside. There are also tables and benches outside for outdoor service, but we quickly learn that the two 'assistant waiters' – the robots – cannot drive out there. Due to their programming, they must stay inside the premises, where guests sit at tables and chairs along the length of the long, narrow room.

In one part of the restaurant facing the street, guests can sit on slightly elevated platforms, a bit isolated from the other guests. In the middle of the restaurant there is an enclosed area with plant beds that are about 1.5 metres high, which partly block the view of the tables inside this area.

We walk over to the bar, which is equipped with draft beer taps and has plenty of spirits on the back wall, as well as fridges with wine, water, and soft drinks.

In the restaurant, customers, both children and adults, human waiters, and the two robot waiters move around among each other. But it is a human waiter who spots new customers and shows them to a table. At the back of the room, we find an entrance to the kitchen and the charging stations for the waiter robots. The customers sitting closest to the kitchen are served by the human waiters, while the robots are more likely to be sent to the tables furthest away. This is an estimation made by the waiters, who judge whether it will be more optimal for them to deliver than to spend time sending the robots. The customers sitting furthest away are served both by waiters or by one of the two robots, which drive back and forth with food on trays built into the robot's 'body'.

The robots are about 1.5 metres tall. They are humanoid which means they include human-like features such as eyes, mouths, arms, or feet and often in the case of human-like robots,

a name. In this case they are white, with round 'heads' that have blue eyes and a mouth on a rectangular display which can both smile and widen and close its eyes. On the back of the robot's body there are a row of shelves where food can be placed. On the front of the 'body' there is a large screen that can display messages, advertisements, campaigns, etc. as for example, presenting the robot as a delivery and marketing expert as well as showing the robot's fictive name (often inspired by science fiction figures from for instance Star Wars). In our experience human staff often give names to humanoid robots. Apart from science fiction inspired names it can also be 'human names' such as Robert. This highlights how people try to make the humanoid robots even more human-like.

On the 'stomach' of these waiter bots we find what looks like two small white, stationary 'nubs' that hold the screen where orders are displayed – they look a bit like hands. Above the torso with the large screen, the small screen in the 'head' (with the robot's eyes and mouth) can change from 'face' to information since the screen shows information about the robot when it is standing still. The little face only appears when it is driving, where it blinks and makes small faces including smiling to guests – maybe to ease its passage through the restaurant. Once again, this example shows how there is an attempt to make robots more human-like, which can also cause robots to seem more autonomous.

It is on the screen in the robot's 'stomach' that the waiters enter which table number it should deliver food to, and on which shelf it is placed. Behind, it has between 3-5 shelves, on which the waiters place the food. When the robot delivers the food, it turns 180 degrees and presents the food to the table. When it turns around, it asks the guests to take the food, and then it drives away by itself.

The robot also 'talks' through a speaker telling the guests about the menu and which tray in the row they can find their food. The sound is not very loud when it delivers food because the speaker points towards the front and it turns its back on costumers, when food is delivered. On the robots 'back' with the rows of trays, there are blue lights on each shelf, which flash and glow when it is driving and indicate which shelf the food should be taken from.

We also notice that the robot's pace is slower than the human waiters due to safety issues.

The robot's tasks and people's interaction with it

Jack, the owner of the restaurant, is a busy man. Even though the robots drive around and help the waiters, he has plenty to do assisting in other ways. The same goes for the other human waiters.

The robots, as mentioned, mainly carry food to the guests sitting furthest away in the room and we wonder about that near the kitchen it is the waiters who deliver the food. When we later make interviews with waiters, they explain that when guests are sitting close to the kitchen, it is quicker for them to walk the food over themselves. It is also the human waiters who clear the dirty plates also in the back of the restaurant. The robot's sole task is to deliver food to the guests further away. As the robots have no functional arms they cannot clean the tables. And furthermore, for food hygiene reasons, dirty plates are not allowed to be put back on the robot's shelves.

Jack explains that the robots are currently not allowed to deliver drinks, as it is not able to handle situations where liquid might spill if it suddenly must stop or hits a threshold or edge. But he expects that it will soon be updated so it can deliver drinks in the future. Both he and the waiters are looking forward to that.

Waiters and robots have certain daily routines in their cooperation. The waiters come out of the restaurant kitchen carrying food in their hands and place the plates and dishes of food on one of the robot's shelves, type the table number the robot should deliver to on the screen on the robot's 'stomach', and place a small sign with the same table number on the shelf.

Then the robot drives off along a pre-programmed route and delivers the food to the guests. After that it drives back and positions itself in front of the kitchen, ready for the next delivery.

We observed that the robot did not drive in straight lines when delivering food, presumably because it tried to drive around people, but this also made it unpredictable for other guests. The same applied when it made small turns and swings while driving, which may have been caused by minor glitches in its programme or small obstacles that it might have scanned.

The waiters have put a small note on the robots saying that guests are not allowed to put dirty plates on the robot. However, we could see that the waiters themselves often ignored this message when they were in the middle of clearing tables. They used the robot's many shelves as a kind of help since it can carry far more plates than they can themselves. So, in practice they overlooked the instructions for how to use the robots and found it helpful in ways that were not stipulated as one of the robot's official functions.

Technical description of the robot

The robot's 'face screen' at rest basically looks like a flat, black tablet. When the eyes appear, they are blue or black, and the mouth can be red. On a handle behind this screen sits the robot's VSLAM sensor.

Visual SLAM (Simultaneous Localisation and Mapping) is a sensor technology that enables the robot to operate based on a map of surroundings while it simultaneously tracks its own position in the environment using visual data from cameras. The cameras work together with a LiDAR unit in the base. LiDAR is a laser scanner that sits around the robot's 'foot' area and produces a 2D 'radar image' of the surroundings. It is mainly used for basic navigation and for detecting obstacles at floor level such as walls, chairs, and tables.

VSLAM, by contrast, is a method that uses camera(s), often directed towards the ceiling and fixed structures, to build up a map of the restaurant and to continuously calculate where the robot is within this map.

Stereo vision is yet another set of cameras, here in pairs, so the robot can perceive depth, like our eyes. They are placed around the robot and are used primarily for close-range safety: detecting low obstacles, edges, stairs, children's legs, bags on the floor, etc., so the robot does not collide or drive off an edge.

In the interaction between these three technologies, LiDAR provides a 2D overview around the base; VSLAM provides the map and precise position in the room; and stereo vision protects the robot from collisions at short range.

The face points in the direction of travel; the trays are on the back. On top, behind the head, there is a red emergency stop button.

The waiter robot in the restaurant functions as an autonomous serving cart that has learned a map of the restaurant with VSLAM. During installation, the robot is driven around so it can build a 3D map of the surroundings and register the kitchen, tables, and preferred routes.

When the staff have placed the food on the robot and selected a table on the touch screen, it uses the map and its sensors to know where it is always and to find a safe route through the restaurant without needing rails or magnetic tape on the floor.

During operation, the robot registers guests, waiters, and other obstacles and slows down or stops until the way is clear, whereupon it continues. At the table, guests are instructed to take the food from the trays, and when trays are empty, the robot drives itself back to the kitchen or on to the next table.

It can also find its own way back to the charging station near the kitchen when it needs power, and it can handle minor changes in the interior layout via its cameras and other sensors.

Despite these technical refinements, the robot is dependent on the maps of the restaurant created via the initial scanning when it was installed, which means the tables cannot be moved around if the robots are to be able to deliver food to them. From the scan it 'knows' where the tables are and what the quickest route to a given table is. It has sensors that register when there is something or someone in front of it, so it can break, stop, and choose a new route.

Twice a year a technician comes and calibrates the scanners. The technician uses the LiDAR sensors to navigate, which work by sending out lasers to make a new map of the robot's surroundings. The sensors register when guests take food from its shelves. It can also play a little melody and communicates through speech (for example, asking to be allowed to pass an obstacle).

It probably has more functions than are used in practice, as there are buttons that can be used to interact with guests in the restaurant, but this was not something we saw in action or that the waiters mentioned.

As in so many other cases, the company behind the robots presents them as more well-functioning and human-like than they are – for example, in a video where a guest pad the robot on its head and the face on the screen blushes and the robot engage in a conversation with guests. In practice we did not see any of these interactions.

When there are several waiter robots in a restaurant, they communicate with each other. This allows them to drive around each other, so they do not collide.

The restaurant's layout

The restaurant has lofty ceilings. There are small tables for two along the walls, and larger tables placed more freely in the room. There are also booths, where you sit on benches facing each other. All these booths are positioned along the outer walls. In these booths, guests sometimes must move their chairs when the robots deliver their food to give it space.

The waiters are careful not to move the tables around too much in the restaurant, not to disturb the robot's orientation. The robots, however, still have more freedom than if the robots were simply running on tracks.

On waiter, Lucy, told us that in a similar restaurant in another city, where she also worked, they had problems with their robots because the room has very long aisles, and sometimes the robots could not figure out where they were and stopped before they were meant to. In the restaurant where she now works, the interior has, in contrast, been arranged so it fits with the robot design.

Training of the restaurant's waiters and guests

Many of the waiters are quite young. The only requirement for being hired is that they must be at least 18 years old. They explain that they do not know much about how the robot's function. They know how to start them and place orders. One evening when we were there, we perceived what we considered an anomaly in the robot's behaviour. We asked a young male waiter about why the robots moved so strangely but he could not say what had happened and just said, "I'm new so… I don't know." It is not only costumers that are new to the robots – also staff can be divided into 'newcomers' and 'experienced' and it makes a difference in how well they understand the robot's behaviour.

The waiters generally described the robots as a good help when they work but also that they often must help the robots, which also underline that breakdowns or malfunctions are normal in the daily routines. The young waiters do not have extensive prior knowledge of the robots or the anomalies that may occur. They are learning by doing. The 'old-timer' tries to teach the new about robot glitches, but they cannot capture everything in words. The experienced waiters explain that they learn how to start the robot when they are hired – but that some of them, over time, develop deeper experience with how to best work around the robots.

Both Jack and Lucy explain that it takes two minutes to train a new waiter to use the robots. We are also shown a typical 'training' by Lucy. She demonstrates what it involves: the waiter learns to pick a tray level, place the food on that shelf, and attach a sign with the table number the food should be delivered to. The table has a pre-programmed number, which is then entered on the robot's 'belly'. After that, the robot is sent off, and the training is done.

Beyond the waiters, who know the robots' routines and limitations, there are the restaurant's guests, who must figure out for themselves in a very short time how to interact with the robots – and this does not always go smoothly. During our fieldwork, we observed how some guests, after taking their food, became irritated that the robot just remained standing at their table once it had delivered the food. It is as if they expected a human behaviour from the robot (no human waiter would linger at a table after the food is delivered) and get annoyed when it's not so human after all. The costumers for instance waved their hands to shoo it away when it did not immediately drive off like you do with an annoying bird. In addition to finding the robots an exciting novelty (and some guests had come precisely to experience the robot waiters), guests also had to adapt their bodily movements to the robot's when they move about, because robots move around at another pace than humans.

The waiters also 'train' costumers on the spot. They always ask whether you have been to the restaurant before, and if you say no, they explain the concept, how to order, and tell you that there is a robot waiter that can deliver the food.

Lucy told us that they have experienced that some older guests may ask not to have their food delivered by the robot. In such cases, the waiter writes a small note with their table number for the kitchen, explaining that the robot should not be sent to them. But of the guests we talked to during our field visits, all said that the robots were fun and easy to use. We for instance talked to young people crowding the restaurant, and a mother with her daughter. Some young men explained that they found the robot less awkward than a human waiter, because the robot does not ask them any questions if they, for example, wave it away. However, they found that during their visit, they themselves became less polite towards the human waiters, as they quickly grew accustomed to not thanking or acknowledging the robot, since it did not require or register such behaviour. The mother said that her daughter thought the robot was fun and therefore liked the restaurant. A young woman described the robot as cute and cosy, but some also said they thought it turns the wrong way when serving the food (because they could not see it's face). This was a common thing, since all the costumers, we spoke to, found it unnatural that the face and speakers are turned away from the guests' table. This is perhaps because the robot's "human-like" appearance, causes the customers to expect it to behave more like a human.

The next sections in the report are about the analysis concepts we have found to be of importance across all cases: Safety, Trust, Anomalies, Sim2Real and Sabotage/Tinkering tied to RoboSAPIENS (see Report V). These themes were identified in Report I as cross-cutting concerns. Report II systematically examines how they manifest specifically in public-facing, mixed-experience environments.

Safety – Restaurant 1

Safety refers to how robots can cause harm to their surroundings. As mentioned, the robots are fairly large, but they move slowly to avoid injuries. Despite this, Lucy explained in an interview that she often experiences them running into her because the sensors don't register her right away. However, they do not cause serious harm, because they do not drive fast and stop when they encounter resistance. She even joked that if this were in the US and not Europe, the restaurant might have been sued, as the robot can occasionally run over a customer's foot. But it sounded as if this is rare, and she was not aware of any permanent injuries to staff or guests.

She emphasised that it can be a problem that many children come to the restaurant, since they are spontaneous and may think the robots can do more than they actually can. When children, for example, run in front of the robots, things can go wrong if the sensors do not register them. Lucy has experienced this herself and says that her own daughter was knocked over by the robot when she was one and a half years old.

According to Jack, it can also happen that the robot runs into a table or something similar, but this does not happen often. They have also taken precautions, he explains, by limiting the robot's opportunity to make a mess: as mentioned it is not allowed to carry drinks, as it is (so far) not stable enough and would spill, especially if it has to brake while delivering.

Jack explains that technically the robots could drive faster, but they are not allowed to because of safety considerations. Most safety problems arise when the robot's sensors are not properly calibrated, or when people walk too quickly in front of the robots so that the robot does not have time to register them.

In our interview with Jack, we also discussed the concept of *perceived safety*, that is, whether people who interact with the robot feel safe and trust it, even if they may not have good reason to. Bodily experience in human-robot interaction plays a role here. We experienced ourselves how the robot can feel unreliable when we tried to stand in front of it. We were unsure whether it had registered us and, if it had, what it would do. Even though we knew it should stop before hitting us, we did not fully trust that it would and stepped back out of its path before it reached us. When we described this experience to the waiters, Lucy said she recognised it – but that, at the beginning, she also believed the robot could sense more than it could, just like her daughter who was hit by it, and that this was probably also why she herself had been hit several times.

Lucy also told us that she does not like having the robots moving around if she is alone in the restaurant in the evening. She clearly does not experience the robot as if it were a human anymore. She says that if the robots have irritated her, she has at times given them a kick. She also described an episode when she felt the robot was following her. She then stopped and considered 'laying it down on the floor' (to deactivate it). She described the robots as 'a silent victim of violence', because it does not say or do anything if she kicks it. This indicates both that waiters like Lucy look for and compare these human-like robots to people when they seek safety, but also that the robots are not able to respond meaningfully. For example, they do not know when they themselves are creating dangerous or uncomfortable situations for people around them.

Trust - Restaurant 1

Trust and safety turn out to be closely linked in our field observations along with a learning curve on the human part in the human-robot interaction, which takes time. Above we heard the story of how Lucy's daughter was knocked over by one of the robots. Lucy's story confirms that one cannot always rely on the robot to stop before hitting someone. In an interview, Lucy elaborated that she sometimes feels nervous about working around the robots. She explained that she returned from maternity leave after the robots had been introduced in the restaurant. Initially she did not trust them at all and was anxious about the idea that the robots might run into her or the guests. Over time she has become more accustomed to them, but her trust is clearly conditional – and is shaped both by subjective experiences (like her daughter's accident) and by visible malfunctions and breakdowns.

At the same time, we also saw that guests often had more patience with the robots than with the human staff. We even experienced that ourselves. When our own order was delayed several times, we did not direct our frustration at the robot but at the human waiter, who seemingly did not understand what had gone wrong. This is an interesting example of how responsibility shifts when technology is perceived as 'not morally accountable'.

Anomalies - Restaurant 1

One evening during our fieldwork, a human waiter dropped some glasses on the floor. She started sweeping them into a pile and had to stop one of the robots from driving straight through the pile by pressing the emergency stop button when it came near. Apparently, the shards were too far down for the robot to detect them, and the sound or sight of breaking glass was not registered by the robot as an 'anomaly'. For the waiters it was different. Lucy immediately came running to help the other waiter sweep the glass away in front of the robot. Meanwhile they had to stop the robot from moving. Only when the path was cleared did Lucy start the robot again, while the other waiter guided it back on track. She kept a hand on the top tray after starting the robot, possibly to be ready to press the red emergency stop button again or to ensure it drove in a straight line and did not turn into the glass shards.

Afterwards she resumed sweeping. But suddenly the other robot came driving towards her and, before she could react, drove straight through the pile of glass shards while her back was turned. She jumped forward when she realised what was happening and pressed the emergency stop button. Lucy then returned and together they tried to remove shards under the robot and sweep away those in front so it could continue. The robot had spread quite a few shards because it had not registered them at all. As it drove, it pushed shards along the floor. One of the shards it shoved in front of its base was quite large and could have caused more serious harm to guests if it had not been removed so quickly by the waiters.

Later, after they had swept away the shards around the last robot and allowed it to move again, it continued pushing a fairly large shard that they had not noticed in front of it. Only when we pointed this out to Lucy did she remove it and afterwards exclaim with an ironic smile, once the robot was back in place, 'The robot will survive' (meaning it was not hurt by the shard). She then went back to help the other waiter remove the remaining glass and asked, "Why doesn't it see it?" as if its humanlike appearance made her expect a more huma-like attention. The other experienced waiter replied: "It doesn't think that far."

Another thing that happened while we were observing was that one of the robots went out to deliver food but returned with the food still on its tray. We asked one of the waiters about this, but he could not say why – he was new. We did not find out if the guests had refused the food and 'sent it back' to the kitchen or the order was mistaken. Afterwards a human waiter took the tray and brought the food out, but she also came back again, so it is possible that it was a wrong order. Here, neither human nor robot notices an anomaly – that the food does not match the table order – until it is sent out.

Sim2Real – Restaurant 1

The robots use a pre-programmed map of the restaurant to drive around and position themselves. The map is not a 'human map' but a machinic and static one consisting of numerical functions. This also means that the robots cannot handle changes to the table layout, which highlights the lack of autonomy the robot has regarding its programming and mapping. Lucy also tells us about the other restaurant she worked at. There, they had problems with robots not registering the tables correctly and ending up stopping at a different table than intended – unlike this restaurant.

In the interview with Jack, he explains:

> *"And like, you know, Google, Google map. First of all, we are making the scan. You know, they have camera. They scan all the restaurants. And we need to put the table down for a particular place, and this mean for example, we cannot change the position for the table."*

To understand his robot waiters, Jack has had to familiarise himself with how they work. He understands that they move according to a map that is programmed when they first receive the robots, and which occasionally has to be recalibrated. They also have LiDAR to scan their surroundings. However, the robots do not scan live in his restaurant; instead, they rely on one scanning that must be recalibrated about twice a year.

In one of the other restaurants, where they used an earlier model of the same robots, one robot had somehow 're-programmed itself' and followed a completely wrong scan of the restaurant, which meant it could not be used. As Jack explains, this is a drawback of the robot not continuously scanning its surroundings and driving according to that.

Even during our own fieldwork, we saw that the robots could have problems with changes in the restaurant layout. One evening the waiters had pushed several tables together to form a long table for a larger group of about 15–20 people. Here the robots became confused by the long table and drove in unnecessarily wide arcs around it, which sometimes created problems both for the waiters and for guests at the other tables. While this might not be an anomalous behaviour from the robot's point of view, it clashes with the restaurant's workflow. It is therefore perceived as an error, despite its programming.

Sabotage and Tinkering – Restaurant 1

Sabotage can be understood in many situated ways, and what is not intended as sabotage by one actor can be experienced as sabotage by another. Across all our cases we have seen many examples of people testing and 'playing' with robots – for example, trying to get them to do something wrong. This is not meant to harm the robot but is a form of what we call tinkering (see Report V, where this aspect is elaborated). It is with this perspective that the following section deals with sabotage of, or tinkering with, the robot's work functions.

A waiter explains that the robots are especially entertaining for children who come to the restaurant, and that this sometimes means that the children 'play' with the robots. This includes stepping in front of them to see if they avoid them, as well as touching them. This form of play is probably not intended as sabotage but expresses people experimenting with and testing the robot. Despite the children's intentions, their actions result in the robot having to stop, find a new route, or being interrupted in its work. Experimenting with technology can also be called tinkering and can be seen as a way of expanding one's practical knowledge about the technology in question.

This fieldwork case is particularly useful for saying something about trust, sabotage/tinkering, and safety. The robots move in a human-dynamic environment. The people they interact with do not necessarily have much experience with robots, so it is important that guests feel they can trust the robots. Additionally, the guests are interacting with the robots for a very limited time, so it is important that they quickly feel a sense of trust in the robots. Perhaps giving the robots, a human-like expression is done to attempt to build this trust. In addition, the robots work up and down alongside humans, who also have to trust them. They must be safe, so they do not pose a safety risk to guests or human waiters.

Restaurant 2 - Waiter Robots accommodating public places to robots

Description of the fieldwork

This restaurant is located in a medium-sized town, at the end of what looks like a pedestrian street. We notice that there are no tables outside the restaurant. To the left of the entrance, the restaurant is furnished with a row of two-person tables along the window facing the pedestrian street, and parallel to this, along a partition wall on the other side, there is also a row of two-person tables. Between these rows there appear to be booth-like seating areas shaped like half-circles. Our main contact here is a waiter named Emma. She meets us on a platform to the right of the entrance, where there are about 6–8 two-person tables. The platform is right next to the entrance to the kitchen, and here the robots stand fenced in by a sort of 'wooden rails'. From here you can see the whole restaurant and straight into their large bar, where Emma is making drinks for some guests. We also get our first clear look at the robots here.

Technical description of the robot – Restaurant 2

The robot at Restaurant 2 looks different from those we saw at Restaurant 1. The biggest difference is that the robot's screen faces the guests when it delivers the food and that its trays extend all the way through its 'body'. In addition, it is less anthropomorphic: it does not have a clear front and back 'body' and therefore does not display messages separately to guests and waiters. But like in Restaurant 1 it is a humanoid. The screen had a face with blue eyes and a mouth that smiled and blinked while it drove around.

This robot does not use VSLAM sensors, so there is no 'marker-free deployment'. It is therefore much less autonomously functioning than at Restaurant 1. Emma tells us that there are several white dots in the ceiling which the robot uses to navigate around the restaurant;

it therefore also lacks VSLAM to perceive the room through cameras. It may run on what we later learn is a RSDI-system that only uses Lidar and sensors (and no camera feedback).

The robot has four shelves, three of which are adjustable, and each shelf can carry a load of up to 10–15 kg, meaning the robot can carry 40–50 kg in total. The robot's navigation and mapping are precisely mapped, and it is those ceiling dots it navigates after.

But the robot also has 3D perception, which provides a form of depth sensing, where technologies such as stereo vision, depth cameras and LiDAR are synthesised to estimate distances and create a three-dimensional representation of an environment. This gives the robot an understanding of an expected room, which it can use to identify "anomalies," such as when a ball comes towards it so that it can react – for example by leaving its planned route briefly and then finding its way back. The robot system furthermore has car-grade suspension where each wheel moves independently for better stability. It is equipped with shock absorbers that reduce vibrations and make the ride smoother, even on uneven surfaces. The design has been developed and optimised using CAE simulations (Computer-Aided Engineering).

Despite these technical refinements, the robot is dependent on the restaurant maps created by scanning the entire restaurant when it was installed, which means the tables cannot be moved if the robots are to deliver to them. From this scan it knows where the tables are and what the quickest route to a given table is. It has sensors that register when there is something or someone in front of it, allowing it to brake and choose a new route.

Every six months a technician comes to calibrate the scanners (we expect because these systems over time change the original calibration and 'grow tired'). The technician uses LiDAR sensors to navigate, which work by sending out lasers that again create a map of the robot's surroundings. The robot also has sensors that detect when guests take food from its trays. It can play a small melody and communicate via speech (for example asking people to move aside so it can pass).

If there are several robot waiters in a restaurant, they communicate with each other so they can drive around and avoid one another. In this sense the robots and waiters together form a 'fleet' communication system in the restaurant. The way it is presented on their website, it indicates that the robots have an advanced form of autonomy, since it can react to the other robots in the fleet and change route, or brake completely to give space. This could potentially also be valuable for the collaboration between robots and humans, if the waiter in the restaurant could become part of the "fleet".

The restaurant's layout and work tasks

To the right in the restaurant room there is a narrow door that frame the entrance into the kitchen and a small ramp over the threshold. This is also an accommodation adjusted to the robots as they cannot otherwise pass the thresholds built into the floors in many places. As in Restaurant 1, the robots stand charging just outside the kitchen, where staff bring out plates and place them on the robots. But at this restaurant the robots also drive directly into the kitchen, where the kitchen staff can put the food on them, so the waiters only must add toppings and set the table number the robot should go to. At one point while talking to

Emma, she says, "Now it's happening," and we can see how one robot drives to the doorway; its interface indicates that it has registered that the other robot is on its way out of the kitchen. In this restaurant the robots are part of the food delivery in the whole restaurant, while the waiters' tasks consist of ensuring that guests have drinks and clearing plates from the tables.

Before our more formal fieldwork visit, we had also visited this place several times. One evening we acted as "guests" ourselves and observed this public space while eating. We saw the robots serving guests and the waiters utilising them. Guests in general showed a high degree of adaptation and patience in their encounters with the robots, stretching physically to reach the food on the robot's trays, especially those seated on the raised platform where the robot cannot get all the way into the table.

This adaptation seems to happen without major dissatisfaction. On the contrary, we observed that guests were more patient with the robots than with the human staff – this also applied to ourselves. Like in Restaurant 1 we ordered food that failed to arrive several times. However, like other guests we did not direct our frustration at the robot, but at the waiter who apparently did not understand what had gone wrong. This is again an interesting example of how responsibility shifts when technology is perceived as 'not co-responsible'. We also observed, like in Restaurant 1, that guests both help and tease the robots, including a guest who deliberately placed food or dishes on the wrong shelves to see how the robot would react – behaviour we have also seen in other contexts with robot technology.

It is worth noting that the robots are not used consistently. Waiters still bring food and collect dishes, even though some guests also place dirty items on the robots' trays. This indicates that the robots do not fully replace the staff, and it does not seem as if they actually reduce the workload significantly in practice. This raises questions about whether the robots' time efficiency outweighs the additional demands of operating and maintaining the system.

Emma told us that the robots are a help when the restaurant is busy, but when it is quiet, they are less helpful, since then the waiters will also spend more time making sure the guests are happy and ask if they need anything.

The next sections in the report are about the cross-cutting concepts: Safety, Trust, Sim2Real, Sabotage/Tinkering and Anomalies tied to RoboSAPIENS (– see Report V).

Safety – Restaurant 2

Emma tells us, like Lucy in Restaurant 1, that she has experienced the robot bumping into her, but she emphasises several times that this only happens if you walk to quickly in front of it. She thinks this is because the robot cannot read such fast movements in time, and therefore she does not consider it the robot's fault. She explains that "it doesn't do it on purpose." She also says that the robot does not hit very hard and that it has never felt violent when she has been hit, because its sensors eventually register that there is something there and the robot stops.

However, the marketing material for the robot says something quite different – that the robot can avoid collisions because its depth perception makes it react quickly. This question of the mismatch between human rhythm and robot rhythm (and the differences between the two kinds of bodies) is a recurring theme, elaborated further in Report IV, which concludes on our robot cases.

According to Emma, the biggest problem they experience with the robots is that they often stop or become 'confused.' As in Restaurant 1, the restaurant has two robots, and during our visit both are in use. Although both restaurants have roughly the same number of guests, this contrasts with Restaurant 1, where we saw that sometimes only one of the robots was used. Emma says the robots coordinate when they have to pass each other. Even though the manufacturer highlights this as a smart feature, it can create problems if people walk by at the same time: the robots may fail to 'see' a person even if they register each other. Emma tells us that one robot once ran into her because it was in the middle of coordinating with the other robot. She thought she could just slip past because they were standing still. But just as she passed, they completed their coordination, and the robot had not 'seen' her in time – and therefore drove into her. This is again a case that show humans, and robotic machines work with different bodily rhythms.

As mentioned earlier, the restaurant has level differences, which means the tables and chairs on the raised platform are not in danger of being hit by the robot. However, Emma tells us that she once stood talking to some customers on the platform and did not hear that the robot had come up behind her. When she stepped down, she walked into the robot. The robot's silence here becomes a safety risk for humans because you do not always realise it is there. In Restaurant 1 the robots play a little when delivering food, to alert guests and waiters of its location, and therefore avoiding such instances.

Emma also explains that if something goes wrong and the robot stops because something blocks it or it runs into something, there is a large red emergency stop button they simply press. When they press the button, the robot stops immediately, and they can push it away from whatever caused it to stop. In the restaurant they have these dots in the ceiling that the robot uses to read its route, and Emma explains that it is important, when pushing the robot back into position, to do so relative to the dots. Otherwise, the robot cannot read where it is and will stop again.

As in Restaurant 1, the costumers and human waiters are not allowed to put used plates and glasses on the robot. When we ask Emma whether this is due to food safety rules – as we were told at Restaurant 1 – she replies that she has been told it is for safety reasons. She says the robots shake quite a lot and that if customers stack dirty dishes on them, there is a substantial risk that they will fall off, which is why they are not allowed to do so. Again, this contradicts the technical descriptions, where the manufacturers claim the robot is equipped with shock absorbers to reduce vibrations. But here the problem is also that the robot cannot meaningfully connect vibrations with the danger of 'dropping' the dishes.

Trust – Restaurant 2

Emma repeatedly tells us that it is not the robot's fault when it hits someone – it is because one walks too quickly past it or because guests put something in its way. She also says that it does not do it "on purpose", and from this, one can infer that she trusts the robot enough to say that it would not willingly cause harm, and that such events are mistakes on her own side.

When we ask whether she trusts the robot, she answers that she trusts the people who made it rather than the robot itself. She says she trusts that it has been programmed properly and that it will not do anything it is not programmed to do. She expands on this by saying she has not experienced it doing anything other than what it is supposed to do – and that is also why she trusts it.

This attitude towards the robot stands in stark contrast to Lucy, who calls their robots a "silent victim of violence" and says that she regularly kicks or hits them. Emma almost apologises on behalf of the robot and says it could not help hitting her – "that's just how it is programmed."

It is her knowledge of the robot's technical background that allows Emma to know it does not intentionally make mistakes or choose to drive into her. Lucy, who has worked at Restaurant 1 longer and has a higher position than Emma, presumably has more experience and knowledge of the robots. Despite (or perhaps because of) this experience, Lucy finds the robots irritating because they bump into things or stop. Emma does not share this irritation but instead seems to feel a kind of sympathy for the robots and their challenges.

Anomaly – Restaurant 2

The tables placed on the platforms led to the emergence of anomalies, since it created a difference in heights which complicated delivering and serving for the robot. In restaurant 2, such as in restaurant 1, a small note had been placed on the robot stating, *"Please do not place anything on the robot."* As mentioned earlier, staff members are also instructed not to transport dirty or empty dishes on the robot. This was explained as a safety measure, as the robot tends to shake, which may cause items to fall off. However, it was also noted that both guests and employees continued to do so. Guests likely do not notice the sign, while employees find that using the robot in this way reduces their workload.

An anomaly in this context can therefore be understood as customers placing empty plates on the robot. Whether this anomaly should be regarded as positive or negative largely depends on the context and the actor involved. If customers place stacked bowls on the robot and these fall off, the consequences are negative, as the items may break. Conversely, if employees occasionally use the robot to place dirty tableware on it during busy periods, this can be seen as a positive anomaly, as it functions as a form of assistance and help for the staff.

Thus, the evaluation of an anomaly as either positive or negative is highly dependent on both context and actor. Additionally, it varies according to the perspective from which the anomaly is viewed. From the restaurant's perspective, it constitutes a negative anomaly, as explicit instructions are given not to place items on the robot. From the manufacturer's perspective, however, it is unclear whether this behaviour should be considered an anomaly at all, or whether it may in fact be an intended or anticipated function.

Sim2Real – Restaurant 2

As described, the robot navigates according to dots in the ceiling and pre-defined routes. From the manufacturer's side, AI algorithms and "multi-robot dispatching" are highlighted as features that can coordinate several robots, so they avoid overlap, waiting time and collisions. In practice, however, this does not seem to work well for Restaurant 2, since the robots' coordination with each other is, according to Emma, the biggest challenge they have. When they communicate about who should go first, they often stop and "think", and in the meantime humans trying to pass can end up being hit – as in Emma's example where she walks past them while they are coordinating and is then hit when they suddenly resume movement.

Sabotage and Tinkering – Restaurant 2

As in Restaurant 1, we see guests experimenting with the robots. Emma describes children who play with them by stepping in front of them and touching them, forcing them to stop or reroute. During our observations we also saw adult guests deliberately put dishes on the wrong shelves to see what would happen. From the robots' perspective, these actions can be seen as "sabotage" of their work functions, but from the human side they can also be seen as tinkering – a playful way of testing and exploring the technology in order to learn more about how it works.

This case from Restaurant 2 shows, among other things, how human-robot interaction can suffer when robot-robot interaction fails, and how humans can end up in a compressed situation when robots must coordinate with each other.

Fieldwork includes examples of robot coordination but also that staff experience that the robot "shuts down" 2–3 times a day, but longer fieldwork could have revealed whether the number is lower or higher over a longer period of time. We are unsure of whether restarting the robot when it stops is really the correct way to handle breakdowns, and how the robots in practice coordinate with each other. Such technical details could contribute to the discussion of positive and negative anomalies. We also do not know whether it is intended or not that the robots should be able to carry dirty dishes – but it seems to be officially prohibited for hygienic reasons even if both guests and staff do it (and robots do not protest).

Receptions and Museums - The HUMANA Robot guiding guests in public spaces

In the following we move on to another kind of robots in public spaces to be found in places like hotels, museums, airports, and places where you have receptions. These types of robots are always humanoids – with human-like features. They also have voices and display where they can talk about what the visitors are to expect from items in the rooms or how to move about in physical premises.

Description of the fieldwork

The specific visit we recount here took place in a large European city, at a medium-sized display site that present new technology including robots to customers. MUSEA is a display site that has invested in a humanoid robot to guide their customers into and around their showroom. They buy diverse kinds of robots, resell them to business clients, and help set them up. Their customers visit to see the newest technology, and the HUMANA robot helps sell these technologies by itself being a piece of innovative technology.

Technical description of the robot – HUMANA

HUMANA is a humanoid robot developed to guide business customers into MUSEA's showroom, where it then introduces various objects in the room. It is programmed to stop at each display case and talk about the newest technologies. The display site itself describes HUMANA as a kind of "wow factor" for businesses, intended to impress their clients.

A technical description of the HUMANA robot: it is roughly the size of a human, with a humanlike face and round eyes that can change colour. When guests arrive at the hotel or reception, it can greet them and welcome them in; on the screen on its torso, guests can read or search for further information about the display site or hotel.

It has arms and hands but no visible legs; instead, what looks like a long white dress is actually the robot's platform, which moves on wheels.

When it has to show guests around, it follows a pre-programmed route, where its own system reads a number of points in its internal map. The route is a so-called "highway" with the points the robot navigates after and tries to move towards. On the back of the robot, as on many other robots, there is a red emergency stop button, which the staff call "the panic button."

It has no camera, but there are sensors on the sides of its "head," in the base, and on the front and the back. HUMANA also has loudspeakers through which it "speaks," with a distinctly robotic artificial voice. What it says is pre-programmed and triggered by several prompts that either appear when you press on the display on its torso (a device like an iPad), or when it reaches certain points on the tour that activate a pre-recorded text. In its face there is also a "mouth," and that is where the loudspeaker is located.

The two eyes appear very lively on this humanoid robot. They follow the people standing in front of it, and it looks after them when they move to the side. HUMANA can also be remote-controlled via an app, but it will still depend on points in the ceiling or floor.

Introduction to Humana onsite

The fieldwork consisted of observations of how the robots were used, conversations and a guided tour with staff from MUSEA, and participation in demonstrations of HUMANA, a robot developed for use in hotels, museums, or as a receptionist in companies.

One of the demonstrations involved HUMANA showing the way into a room where it was supposed to describe and talk about selected objects in the space (for example, a museum robot might talk about paintings on the walls, a hotel robot might show guests the rooms, and a receptionist robot in a company might show visitors around showrooms where, for instance, technologies are exhibited that need to be explained).

HUMANA began by welcoming us to the tour and then started moving down a hallway towards the room where the objects it was to explain were exhibited. The showroom was at the end of the hallway, in a room with an open glass door set in a large glass wall. Inside the room you could see a number of display cases with objects and a table in front of an information board.

The robot's pace is extremely slow as it moves down the hallway, and Jonas, who runs the place, explains that this is for safety reasons. But it is also an example of how the bodily rhythms of humans and robots do not always match. People visiting a museum, hotel or reception will typically walk faster and move more agilely around the robot, which Jonas explains can confuse it. In this case, the anthropologist moves into the robot's path so that it stops and makes sure to avoid her. Then it continues and resumes the route.

But instead of going into the showroom through the open glass door, it now goes about a metre too far and positions itself in front of the glass wall, saying: "Please open the door so I can go through." It is not 'aware' that it is standing in front of a glass wall (and not the door) – or that it has just passed an open door. This is because, for the robot, there is nothing meaningful in the room; its 'autonomy' is tied to the positional points it has been pre-programmed with. So, if it 'miscounts' numerically, it will also miscalculate the route and therefore assume it is in front of a closed door rather than a glass wall.

In other words, the robot is blind to everything that makes the room meaningful to humans. Although it looks humanlike, it is not able to understand itself in relation to its surroundings. Nor does it understand how to search for novel solutions (such as moving a bit forward or back to search for a door opening). In this case Jonas must press the red emergency button. HUMANA's head drops, and it looks as if it is sleeping. In this state Jonas can roll it back to the starting position at the entrance and start the tour over.

The next sections in the report are about the analysis concepts (Safety, Trust, Anomalies, Sim2Real and Sabotage/Tinkering tied to RoboSAPIENS – see Report V).

Safety - MUSEA

At MUSEA, safety was ensured through low speeds and emergency stop functions. The HUMANA robot moves slowly because it is intended to move around in a room with many people. The slow speed allows it to stop or adjust its direction to avoid running into people. It is very cautious when it moves around and asks for space before it starts guiding.

HUMANA for instance says: “I'm going to start moving to the destination. Please, give me some space to move.” These communications are especially important for visitors, which get a warning every time the robots make a new move.

The display site owner Jonas also explained that when they tested HUMANA with employees from MUSEA acting as customers, he asked the staff to stand still in ways that could not always be demanded of museum guests, who would have to adjust their movements to those of the robots. He did this because, as he says, “I have to tell this to visitors – they cannot just move about. Because the robot… would move and… and see people and then move here and there.” This would make the interactions between visitors and robot unstable. Because HUMANA is programmed to avoid people, it can be confused if it sees people moving around, as it will then try to avoid them. Therefore, it is best to stand behind HUMANA and preferably stand still until it has passed. In this case people are asked to adjust to the robot’s behaviour, but this cannot always be the case with outside visitors. When they have visitors that meet the robot for the first time, they often call the company that sold it to them for maintenance and check-ups and run several tests before they let ‘newcomers’ in. This speaks to the autonomy of the robot, since it shows that they must test it thoroughly before showcasing it, to make sure it works.

Trust - MUSEA

At MUSEA, HUMANA was used to create a “wow” effect rather than to deliver real functionality, which indicates a superficial form of trust. It may create more trust in MUSEA as a credible place in the eyes of their customers – when it works. The display site likes the robot because it shows them as ‘future-oriented,’ but they get worried and begin to mistrust the robot if it repeatedly fails. That is why they so often call the company that sold it to them, because they have learned over time that they cannot feel sure the robot will behave as expected. It thus creates a form of distrust when the robot repeatedly does not work as expected. This is not a mistrust in safety, but rather as mistrust the robot will not represent MUSEA in the best way. The display site owner Jonas said that he was not nervous that HUMANA would drive wrongly or run into anyone because it moves so slowly. For him, his trust in HUMANA capability not to cause harm is directly linked to its slow speed. This is also interesting in relation to the fact that the display site wants the robot to look as smooth as possible, but at the same time they do not want it to move much faster, as that would affect their trust in it and its safety.

Anomalies - MUSEA

HUMANA was often misled by its surroundings, for example by a glass wall, as in the situation mentioned above. Here HUMANA started to say: “Please open the door so I can go through,” which HUMANA repeated until Jonas pressed the red panic button. He did not know why it did this but could say that it is not unusual for the robot to misread the points in the floor or ceiling.

Another anomaly occurred on other tours, where HUMANA managed to reach the room it was supposed to reach (for example a museum room), but then stopped and began explaining a painting that was on the opposite wall, or in another case talked about the different display cases even though it was not standing in front of them

Sim2Real - MUSEA

HUMANA follows a pre-programmed route. The display site owner Jonas describes it like this: “She has a pretty fine path, they call it a highway, which is a set of points that she tries to follow. That's the destination.” HUMANA tries to follow these points, and if a person or another obstacle stands in the way, it will attempt to drive around it and then find its way back to its pre-programmed route. This can lead to problems, which we will return to later.

Sabotage and Tinkering - MUSEA

An example of something perceived as sabotage from the robot’s point of view could be when the visitors step in front of HUMANA. This was also done by the anthropologists as it began its guided tour because the robot’s pace was too slow for her human gait. This, Jonas confirms, might also happen if there is a group of customers HUMANA is supposed to guide. Here the anthropologists were also going in front of the HUMANA robot to see what would happen and how it would react. Jonas could not explain the technical details of what might go wrong, but something definitely went wrong as HUMANA ended up in front of a glass wall instead of going through the door it was supposed to use. This mistake caused the very communicative HUMANA to stand there repeating: “Please open the door so I can go through.” This meant Jonas had to restart the presentation from the beginning and he had experienced this before, and was a bit annoyed he could not just move the robot to the door. Instead, he had to wheel it to the starting position at the reception and start all over. Though costumers at a real customer visit do not step in front on purpose, Jonas worry if the robot could actually ‘sabotage’ the image of the MUSEA company as a whole when it is not working as expected and thus damage the impression of all the other robots and technologies they are trying to sell. However, he also confirms that when HUMANA works it is a real selling point for MUSEA as visitors see it as an almost human all-knowledgeable guide.

HUMANA’s human-like appearance and communicative style amplify this mismatch: although it looks and speaks like an intelligent guide, its autonomy is limited to following pre-programmed points. The more convincingly human the robot appears, the more visible—and disruptive—its lack of situational understanding becomes when something goes wrong

This is exactly the “wow effect” that they want HUMANA to give, and it can be argued that MUSEA oversells what their robots can do when they program HUMANA to appear more “intelligent” than it is and that this can become a problem for the robot engineers themselves. This is very similar to what we have experienced at conferences, where various robots are presented as if they are much more self-thinking and autonomous than they actually are. We have also heard from several robot manufacturers that this can lead customers to believe that robots can do much more than they really can (like argued in Report I).

Shops and Cleaning Robots - The MEASURE Robot meeting people in public spaces

Description of the fieldwork

We have visited several semi-public spaces where robots move around people in daytime of nighttime to perform chores and tasks that are not directly connected to communicating with costumers or visitors. These could be cleaning robots in educational settings and hospitals, or stocktaking robots in shops. In the case described here the visits took place in a European metropolitan city. In companies with robots (see also Report III) we find that in some cases the robots move into the companies' public spaces. In the case we describe here, the anthropologists visited a company with a robot that was not there for the costumer's sake but still encountered costumers in the public space of the shop. SHOPFLOOR is a large chain of clothing stores. In several of the countries where they operate, they are experimenting with using robots to serve customers or as in this case to monitor goods in their shops. MEASSURE is a robot that operates in several of their stores. It is only used for stocktaking and inventory control, not for communication with costumers. The particular store we visited is located about an hour's drive outside the city, in a shopping centre. The centre is in an industrial area and is designed on the assumption that people arrive by car.

The robot was in the stockroom when we met it for the first time. It was then set into motion through the store while there were customers inside in the morning. We followed it around the shop and observed the reactions of staff and customers. Here we could observe a clear difference between the "newly arrived" customers (who perhaps had never seen an autonomous robot before) and the staff, who had learned to interact with the robot over time.

We spoke with an older employee in the shop, who was the one who programmed the robot to drive to its charging station in the morning. She had never been afraid of the robot. The same was the case for a younger employee we spoke to. The staff seemed completely comfortable walking in front of the robot and were not nervous that it might bump into them. By contrast, several customers made a wide detour when the robot came driving towards them. Some customers seemed almost a bit frightened or angry about the robot's presence, while others looked at it with curiosity. Contrary to other types of service robots in the wild, these robots did not have a human-like face, which might have eased the contact with costumers.

The robot - MEASURE

MEASSURE is a functional, non-humanoid robot for stock measurement and inventory control. It is a tall white box with various built-in measuring instruments. The robot drives around the store once during the day when there are customers, but primarily at night.

We came to observe how it moved through the store as it returns to its charging station in the morning after a night's work.

The clothing chain's stores typically consist of long aisles with diverse types of garments (for example shelves with caps or T-shirts). All clothing has a tag, a label, which can be read by the robot. During the night, the robot drives around and scans how much clothing has disappeared from stands and shelves, and informs the staff if items need to be restocked. At night, when the shop is closed, it drives around and scans all the items that have not yet been sold. It then compares this count with the stock figures it receives from the lorries delivering goods and with the staff's records of the day's sales. This allows staff to detect theft from the store if the two counts do not match.

The robot ends its tour in the stockroom, where it stands overnight. In the morning it is again set to drive through the store and over to its charging station. The staff affectionately give the robot nicknames inspired by iconic robots from Star Wars and generally seem completely comfortable with its presence.

Like many other robots in public spaces, the robot has a predefined route, which it follows using a number of reference points, but it can deviate from this when it encounters employees or customers who step in front of it. It can both stop and drive around them and then find its own way back to its route afterwards.

Technical description of the robot

MEASSURE is an autonomous inventory robot with RFID-based registration of goods. It is about 2 metres tall and has sensors at the front and the back. Unlike many other robots in public spaces, this is not a humanoid, but simply a moving "column" that does not try to appeal through a human-like appearance. The sensors that scan the clothing tags are not visible from the outside.

The robot has a tiny radio scanner, and the room has lots of small stickers (RFID tags) hidden around: each sticker sends a short ID and an indication of how strong the signal is when the robot listens. As the robot moves, it reads tag IDs and signal strengths, which tell it roughly where it is, and which objects or places are nearby. The robot filters out noisy readings, uses patterns of signal strengths and tag locations it has learned earlier (a map), and matches current reads to that map to figure out its position.

To follow a route and avoid obstacles, the robot combines the RFID-based position estimate with its other sensors (such as cameras or bump sensors). The map provides a planned path;

RFID helps confirm where the robot is along that path. If RFID or other sensors detect something blocking the way, the robot recalculates a short detour and uses updated RFID reads plus its sensors to stay on course and avoid collisions. In short: RFID provides location clues and object IDs, the robot matches those clues to its map, and it adjusts its movement in real time to follow the route and dodge obstacles.

RSDI, VSLAM, and CAE are three different classes of techniques used in engineering and robotics with different goals, inputs, and outputs:

- **RSDI (range sensor–driven inference):** Uses distance/depth sensors (LiDAR, sonar, depth cameras) to infer robot pose and build geometric maps of the environment.
- **VSLAM (Visual SLAM):** Uses camera images to simultaneously estimate camera/robot motion and build a map with visual landmarks (appearance + geometry).
- **CAE (Computer-Aided Engineering) simulations:** Use physics-based numerical models (finite element analysis, CFD, multibody dynamics) to predict how engineered designs behave under loads, thermal conditions, fluid flow, etc.

The Visual SLAM is the technique that may open for most questions of surveillance in public spaces such as the ones presented here ranging from restaurants, museums, hotels, and shops.

The next sections in the report are about the analysis concepts (Safety, Trust, Anomalies, Sim2Real and Sabotage/Tinkering tied to RoboSAPIENS – see Report V).

Safety -SHOPFLOOR

At SHOPFLOOR, MEASSURE was primarily used outside opening hours, which minimised the risk to customers and staff. It was only in the morning, when the staff sends MEASSURE to its charging station, that it drives through the store. Customers are otherwise not exposed to it. It drives very slowly, which the staff explained as: "And it has to go slow for security."

Trust -SHOPFLOOR

MEASSURE appeared to be accepted by employees as a useful tool and almost just another colleague in the shop, whereas customers showed scepticism and distrust when encountering the robot. A concrete example was a customer who, during the field visit, saw MEASSURE and flipped it off. When the anthropologist asked him about it, he replied that he did not like robots and hurried on. There is a clear difference between the customers, who gave the robot somewhat sceptical looks, and the employees, who simply ignored it.

Possibly this is because the employees are used to it. They have learned what to expect so they know that nothing will happen, and that it will avoid them and so on, while customers may be seeing it for the first time and therefore do not know what to expect.

It is also interesting that, if MEASSURE is to be used to detect theft, it then assumes an authoritarian role vis-à-vis the staff. This requires MEASSURE to be precise and not make mistakes, which, according to the manufacturer, it can in fact do (make mistakes): "We have sometimes problems because it takes up this [issue of missing clothes]. And then we have tags from other places, and we detect extra tags sometimes." They have had problems where MEASSURE occasionally counts incorrectly because it picks up the wrong signals, which means it would not be accurate enough for the shop to rely on its counts.

When employees become aware of this, it could lead to mistrust of the robot's work and distrust of management if they use the robot's counts to assess whether there is a theft problem in the store. However, the employees generally trust the robot's physical presence. They are not afraid that it will fall or hit them: "If there is an obstacle it will go around," as one employee puts it.

Anomalies -SHOPFLOOR

One anomaly could be when MEASSURE picks up the wrong tags and thus produces an incorrect count. This happens because there is an error in its internal systems, where it reads tags from other floors or shops. The staff do not know exactly how this happens.

Sim2Real -SHOPFLOOR

MEASSURE operated according to simple routes based on RFID data. It has a map of the store that it follows and that can be used to find out where different goods are located, because it registers them on the map when it does its counts. In addition, like the other public-facing robots in this Report II, it has a predefined route that it follows.

Sabotage and Tinkering -SHOPFLOOR

Interestingly in this case we did not see the costumers, who encountered the robot on its morning route tinkering or trying to sabotage it. They rather seem to avoid it and in one case (as mentioned) flipped it off. We wonder if this may be due to its lack of humanoid features, which would have made it more human-like, but maybe also raised expectations of its behaviour. However, it raises the question: Are humanoid robots especially exposed to tinkering (playing with the robot) or direct human intervention to harm the robot (sabotage)? Furthermore, does a humanoid design make the robot less or safer, both regarding perceived safety and trust, but also tinkering and sabotage? You can't look it in the eyes, so it might be harder to see where it goes.

Understanding Service Robots in Public Spaces

What makes these cases distinct is not only that the robots operate “among people,” but that they operate in publicness: a social condition characterised by constant turnover, low shared routines, and brief encounters. Staff may become experienced, but guests, customers, and visitors remain newcomers. This shapes safety and trust (which must be negotiated instantly), increases tinkering (testing and play), and makes Sim2Real fragilities visible whenever the environment departs even slightly from the robot’s assumptions (crowds, thresholds, rearranged tables, glass walls). Publicness thus reveals how autonomy is not simply a technical property but a situated achievement that relies on ongoing human adjustment.

Across the cases in this report, humanoid design emerges as a recurring source of tension rather than a solution. While faces, voices, and human-like bodies are often intended to make robots appear more approachable and trustworthy, in practice they raise expectations the robots cannot meet. Guests expect human-like timing, awareness, and responsiveness; when these expectations are disappointed, irritation, tinkering, and even physical aggression toward the robots follow. At the same time, humanoid appearance does not lead people to treat robots as morally accountable actors: responsibility for breakdowns is consistently displaced onto humans, not attributed to the robot itself. Rather than bridging the gap between technical autonomy and social interaction, humanoid design often widens it by making the robot appear more capable, intentional, and self-directed than it actually is.

During fieldwork of robots in public places, we encountered various kinds of system where the robots navigate public and more private spaces, and across cases there seems to be a lot of learning going on when robots are introduced in public spaces. Here we find what we can distinguish as ‘experienced’ and ‘newcomers.’ The permanent staff eventually become experienced, whereas the costumers, guests and visitors are all newcomers that have to learn fast what to expect. Sometimes the human-robot interaction goes wrong for both staff and newcomers – often tied to different expectations of safety and trust.

Emma’s attitude to the robot is also interesting: she repeatedly “apologises” on its behalf and insists that it does not hit on purpose, but that she herself has been too inattentive or optimistic. Perhaps because she is “just” a waiter, and not a boss, whereas Lucy might feel more comfortable with complaining and blaming the robot. Besides this can also be seen as an example of *social desirability bias* where informants will answer in a way that is construed as desirable by others. Here Emma might feel that she must present the robots in a positive way, because she might think this is what her workplace might prefer.

Many of the robots in public places have humanoid features and that is both good and bad from a roboticists point of view. Like in Report I this may be a feature ‘overselling’ the robots capacities making costumers expect at more humanlike behaviour than what they get. The human feature for instance cause expectation in robot behaviour regarding the rhythm and pace they move around with as well as the robot’s behaviours when they have a glitch (cannot find a door, push glass shards around with them, or stand too long at a table).

On the other hand, we have not seen (neither the temporary visitors but especially among the more experienced staff), that humans treat the humanoid like they would treat humans. They do, for instance, not blame them for accidents and they may kick them and tinker with them in ways they would never do to humans.

Jack told us that he has read about waiter robots being used everywhere in China. He gave a concrete example: if you order room service on the 27th floor of a hotel, a robot rather than a human comes to deliver the food. He also thought (presumably because he is the restaurant owner) that everyone will use waiter robots in the future, because they do not call in sick, they do not complain, and they need fewer breaks. Lucy also thought that robots will appear everywhere and probably become more widespread, but she added that she herself prefers human service.

SHOPFLOOR's employees also emphasised the need for balance between humans and robots in the work environment. One employee stated that she preferred there to be a balance between the number of robots and humans in a shop. When the anthropologist asked concretely whether she would become lonely if there were no other people in the store, she replied:

"Yeah. A lot. I'm wondering what happens [in the future]... What do I think about the robots and the relationship that there are between people and robots... Hmm?"

It is also interesting to discuss what it would take to improve the technical disadvantages of the robots, for example with more cameras. If the robot had cameras that could better distinguish between people and other objects in front of it, it creates an ethical problem. If such cameras move around filming the robot's surroundings in order to differentiate between people and objects – and perhaps even adults and children, waiters, and guests – this raises problems, since it would mean that all guests are under surveillance. These are people who may not be aware that they are being recorded or whether the robot stores the recordings for a period. T

Additionally, reflections on ethics in HUMANA were also linked to data use, surveillance, and social integration and not so much its humanoid appearance. MUSEA pointed to data handling as the greatest risk. Waiter robots and social robots such as HUMANA raise questions about the future of service work, where machines can change professional identity and interactions.

Conclusion

Public robots are often more human-like than robots found in industrial or dynamic settings (see Reports III and IV). Across sites we find that he more human-like the robot looks the higher the expectation of autonomy. Report I argued that autonomy is often overstated and that real-world functioning depends on situated coordination. The public settings in Report II show why: publicness forces robots to operate with newcomers, limited time for learning, and little tolerance for friction. This is where overselling becomes immediately visible and where Sim2Real gaps appear as safety concerns, delays, confusion, and the need for human intervention. In other words, publicness makes clear that what looks like autonomous operation is often sustained by invisible human work—staff guidance, guest adaptation, and continual repair in practice.

The robots we have seen in action in these cases are generally well-functioning, but it is clear that in many ways humans have to adapt more to the robots than the robots adapt to humans. This applies to both staff and guests. Even though some of the robots have human-like features, their rhythm, flexibility, and movements are often very different from humans. They are able to "understand" where they are in a room, but this position is not meaningful to them. Therefore, they may stop or drive straight towards people, who then have to move quickly out of the way. They are a help to the staff, but can also be a source of irritation, for example, if they do not move quickly enough, or get in the way. For the robot, the world is not meaningful – it is just numbers and calculations. People constantly tinker with but also adapt to the robots; robots help people, but people also help the robots. And the staff (just as costumers) do not believe that responsibility can be attributed to the robots when something goes wrong. It is solely the humans' fault. Even if the robots are humanoid, and this create some expectations of behaviours, robots are also considered different from humans especially then it comes to placing responsibility.

One of the clearest cross-cutting findings in this report concerns humanoid design. While anthropomorphic features are often intended to increase trust and ease interaction, our cases show that they frequently have the opposite effect. Humanoid robots raise expectations of awareness, intention, and responsiveness that current systems cannot meet, while still not being treated as morally responsible actors when failures occur. In this sense, humanoid design does not resolve the problem of autonomy identified in Report I; it intensifies it by overselling capabilities and making the autonomy gap more visible in everyday interaction.

Overall, these are some of the main findings in relation to what makes robots in public places special. In the subsequent reports (Report III and IV) we concentrate on robots in other types of environments. In Report III we mainly visit companies that have involved robots and co-bots in their daily routines and in Report IV we visit robots that work in unusually volatile and dynamic environments such as the ocean or an agricultural field. Finally, in Report V, we shall discuss all the findings across the case-studies in relation to the identified issues of safety, trust, Sim2Real, sabotage/tinkering and anomalies– as well as the ethical questions raised across these studies.

Acknowledgements

The work presented here is supported by the RoboSAPIENS project funded by the European Commission's Horizon Europe programme under grant agreement number 101133807.

About the authors:

Cathrine Hasse is Professor of Learning and Technology at DPU, Aarhus University, Denmark, and research program leader of *Future Technologies, Culture and Learning*. She is also Honorary Professor of Techno-Anthropology at Aalborg University. Trained as a cultural anthropologist, her research focuses on learning, technology, and human–robot relations in real-world settings. She has led and coordinated numerous interdisciplinary research projects, including Horizon 2020 projects, working closely with researchers from engineering, psychology, and the natural sciences. She is the author of *An Anthropology of Learning* (Springer) and *Posthumanist Learning* (Routledge), and has published extensively on robots, learning, and technology in everyday life.

Pernille Maja Carlsen holds a bachelor's degree in Anthropology from Aarhus University. She is currently pursuing a master's degree in IT, Communication, and Organization at Aarhus School of Business and Social Sciences, where she is writing her master's thesis on hybrid work and digital communication. RoboSAPIENS is her first research project, and her interest in the field stems from earlier studies of robotics and technology. She is currently working as a research assistant and has been part of the project since the summer of 2024.

Peter Hommel Østerlund is a master's student of Anthropology at Aarhus University, from where he also got his bachelor's degree in Anthropology. He passed his bachelor's degree with a study on prosthetics and how people with prosthetics are affected by them in their daily lives. He is a student assistant on the research project, RoboSAPIENS, where he works closely with Cathrine Hasse and Pernille Maja Carlsen.